

IBM Cloud App Management  
Version 2019 Release 4

*User's Guide*



**Note**

Before using this information and the product it supports, read the information in [Chapter 6, “Notices,”](#) on [page 65](#).

This edition applies to version 2019.4.0 of IBM® Cloud App Management and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2018, 2019.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>Chapter 1. What's new.....</b>	<b>1</b>
<b>Chapter 2. Known issues and limitations.....</b>	<b>17</b>
<b>Chapter 3. PDF documentation.....</b>	<b>19</b>
Documentation.....	19
Conventions used in the documentation.....	19
<b>Chapter 4. Product overview.....</b>	<b>21</b>
Offerings and features.....	23
User interface.....	26
Getting started: Manage dynamic application and infrastructure environments .....	27
Getting started: Collaborate to rapidly resolve problems .....	30
Getting Started: Proactively manage the health of your application environment – regardless of size .....	34
Getting started: Accelerate your transition to the cloud with DevOps.....	38
Getting started: Performing SRE functions.....	43
Incident Resolution Flow for a Kubernetes App.....	45
Incident Resolution Flow for a traditional VM-based App.....	48
Agents, data collectors, and plug-ins.....	49
Change history.....	52
Descriptions.....	54
<b>Chapter 5. Accessibility features.....</b>	<b>63</b>
Trademarks.....	66
Terms and conditions for product documentation.....	67
IBM Online Privacy Statement.....	67
<b>Chapter 7. Planning your deployment.....</b>	<b>69</b>
Product components.....	70
Part numbers.....	71
System requirements.....	75
Planning hardware and sizing .....	77
Planning ports.....	80
Scaling stateless and stateful services .....	83
Scaling Cassandra .....	84
Scaling up CouchDB.....	85
Scaling up Kafka brokers.....	86
Scaling up ZooKeeper.....	87
Scaling up Datalayer.....	87
<b>Chapter 8. Deploying IBM Cloud Private.....</b>	<b>89</b>
Preparing to install IBM Cloud Private.....	89
Installing an IBM Cloud Private Enterprise environment.....	90
IBM Cloud Private postinstallation manual tasks.....	92
<b>Chapter 9. Installing IBM Cloud App Management - the options.....</b>	<b>93</b>
Installing IBM Cloud App Management with IBM Cloud Pak for Multicloud Management.....	94
Onboarding LDAP users.....	94

Online.....	96
Offline.....	109
Installing the ICAM klusterlet on the managed cluster with Helm.....	133
Uninstalling the ICAM klusterlet with Helm.....	135
Installing the ICAM klusterlet on the managed cluster without helm.....	135
Uninstalling the ICAM klusterlet without Helm.....	139
Deploying agents and data collectors for IBM Cloud Pak for Multicloud Management.....	140
Offline: Installing IBM Cloud App Management stand-alone on Red Hat OpenShift.....	141
Offline: Installing IBM Cloud App Management on IBM Cloud Private .....	148
Planning for a high availability installation.....	148
Certificates for HTTPS communications.....	152
Optimizing performance.....	156
Installing IBM Cloud App Management stand-alone on IBM Cloud Private.....	163
Creating your service instance.....	174
Starting the Cloud App Management UI.....	176
Backing up and restoring .....	177
Back up Helm release.....	179
<b>Chapter 10. Configuring IBM Cloud App Management.....</b>	<b>181</b>
Upgrading IBM Cloud App Management with IBM Cloud Pak for Multicloud Management from eventing only to full monitoring mode.....	181
<b>Chapter 11. Uninstalling IBM Cloud App Management .....</b>	<b>185</b>
Deleting the IBM Cloud Private service instance.....	185
Uninstalling IBM Cloud App Management.....	185
<b>Chapter 12. Upgrading IBM Cloud App Management.....</b>	<b>189</b>
<b>Chapter 13. Deploying ICAM Agents.....</b>	<b>193</b>
Planning agent deployment.....	193
Downloading agents and data collectors from Passport Advantage.....	194
Configuring the downloaded images.....	194
Installing agents on UNIX systems.....	196
Preinstallation on AIX systems.....	196
Preinstallation on Solaris systems.....	197
Installing agents.....	198
Installing agents on Linux systems.....	200
Preinstallation on Linux systems.....	201
Installing agents.....	202
Installing agents on Windows systems.....	204
Preinstallation on Windows systems.....	205
Installing agents.....	206
Installing agents silently.....	208
Installing agents as a non-root user.....	209
Postinstallation tasks for the agents.....	210
Securing the agent installation files.....	212
Bypassing the prerequisite scanner.....	213
Uninstalling your agents.....	214
WebSphere Applications agent: Unconfiguring the data collector.....	216
<b>Chapter 14. Configuring your agents.....</b>	<b>225</b>
Common procedures.....	225
Using agent commands.....	226
Starting agents as a non-root user.....	230
Configuring agents as a non-root user.....	231
Disabling automatic agent start on AIX and Linux systems.....	232
Configuring Amazon EC2 monitoring.....	232

Configuration parameters for the Amazon EC2 agent.....	233
Configuring the agent on Windows systems.....	235
Configuring the agent by responding to prompts.....	236
Configuring the agent by using the silent response file.....	237
Configuring AWS Elastic Load Balancer monitoring.....	238
Configuration parameters for the Amazon ELB agent.....	238
Configuring the agent on Windows systems.....	239
Configuring the agent by responding to prompts.....	239
Configuring the agent by using the silent response file.....	241
Configuring Azure Compute monitoring.....	242
Azure Compute Configuration Information.....	243
Configuration parameters for the Azure Compute agent.....	244
Configuring the agent on Windows systems.....	246
Configuring the agent by responding to prompts.....	247
Configuring the agent by using the silent response file.....	248
Configuring Cisco Unified Computing System (UCS) monitoring.....	250
Creating a user and granting required permissions.....	251
Configuring the agent on Windows systems.....	251
Configuring the agent by using the silent response file.....	252
Configuring the agent on Linux systems.....	253
Configuration parameters for the agent.....	254
Configuration parameters for the data provider.....	255
Enabling SSL communication with Cisco UCS data sources.....	255
Increasing the Java heap size.....	256
Configuring Cassandra monitoring.....	257
Configuring the agent on Windows systems.....	257
Configuring the agent on Linux systems.....	258
Configuring the agent by using the silent response file.....	258
Configuration parameters of the agent.....	259
Configuring Citrix Virtual Desktop Infrastructure monitoring.....	259
Configuration parameters for the Citrix VDI agent.....	261
Configuring the agent on Windows systems.....	262
Configuring the agent by responding to prompts.....	263
Configuring the agent by using the silent response file.....	264
Enabling Citrix read-only administrator privileges.....	265
Enabling monitoring of Windows events and PowerShell metrics.....	266
Configuring CouchDB monitoring.....	266
Configuring the agent on Windows systems.....	267
Configuring the agent on Linux systems.....	268
Configuring the agent by using the silent response file.....	268
Configuration parameters for the agent.....	269
Configuring WebLogic monitoring.....	270
Configuration parameters for the WebLogic agent.....	272
Configuring the agent on Windows systems.....	272
Configuring the agent by responding to prompts.....	276
Configuring the agent by using the silent response file.....	276
Configuring DataPower monitoring.....	278
Configuring DataPower appliances.....	278
Configuring the DataPower agent.....	279
Configuring Db2 monitoring .....	282
Configuring the agent on Windows systems.....	283
Configuring the agent on Linux or UNIX systems.....	284
Configuring the agent by using the silent response file.....	286
Granting privileges for viewing Db2 agent metrics.....	287
Configuring local environment variables.....	288
Prerequisites for Remote Monitoring.....	291
Configuring Hadoop monitoring.....	292
Configuring the agent on Windows systems .....	293

Configuring the agent on Linux and AIX systems.....	296
Configuring the agent by using the silent response file.....	299
Granting permission to non-admin users.....	300
Configuring the dashboard for viewing Hadoop events.....	300
Configuring HTTP Server agent monitoring.....	301
HTTP Server agent code samples.....	303
<b>Configuring IBM Integration Bus monitoring.....</b>	<b>304</b>
Configuring the IBM Integration Bus agent.....	305
Configuring IBM Integration Bus for data enablement.....	308
Disabling snapshot data collection for the agent.....	312
Configuring InfoSphere DataStage monitoring.....	313
Configuring the agent on Windows systems.....	313
Configuring the agent on Linux systems.....	313
Configuring environment variables.....	314
Configuring the agent by using the silent response file.....	314
Configuration parameters of the agent.....	315
Configuring JBoss monitoring.....	316
Add a JBoss Server Management User.....	317
Enable JMX MBean Server Connections.....	318
Enabling Web/HTTP Statistic Collection.....	319
Configuration Parameters for the JBoss agent.....	321
Configuring the agent on Windows systems.....	322
Configuring the agent by responding to prompts.....	322
Configuring the agent by using the silent response file.....	323
Configuring Linux KVM monitoring.....	324
Creating a user and granting required permissions.....	325
Configuring protocols.....	325
Configuring a connection to the RHEVM server.....	329
Configuring a connection to the RHEVH server.....	331
Configuration parameters to connect to the RHEVM server.....	331
Configuration parameters to connect to the RHEVH server.....	333
Configuring MariaDB monitoring.....	335
Configuring the agent on Windows systems.....	335
Configuring the agent on Linux systems.....	336
Configuring the agent by using the silent response file.....	337
Configuring Microsoft .NET monitoring.....	338
Permissions to run an agent by using a local or domain account.....	338
Configuring the agent on Windows systems.....	339
Changing the user account.....	339
Configuring Microsoft Active Directory monitoring.....	340
Running the Microsoft Active Directory agent as an administrator user.....	340
Configuring local environment variables.....	341
Running Microsoft Active Directory agent as a non-administrator user.....	342
Configuring Microsoft Cluster Server monitoring.....	344
Creating a generic service cluster resource on Windows Server 2008, 2012, 2016, and 2019 systems .....	345
Configuring the agent by using the silent response file.....	345
Changing the user account.....	346
Configuring Microsoft Exchange Server monitoring.....	346
Creating users.....	347
Assigning administrator rights to the Exchange Server user.....	349
Making the Exchange Server user a local administrator.....	350
Configuring the Exchange Server for reachability.....	351
Configuring the agent to run under the domain user.....	352
Configuring the agent locally.....	353
Configuring the agent by using the silent response file.....	357
Configuring local environment variables for the agent.....	357
Configuring Microsoft Hyper-V monitoring.....	358

Providing Local Security Policy for running Monitoring Agent for Microsoft Hyper-V Server on Windows by a Non-Administrator user.....	359
Granting Local Security Policy permissions.....	359
Modifying DCOM permissions.....	360
Adding a non-administrator user in the Hyper-V administrator users group.....	361
Adding a non-administrator user in the Performance Monitor users group.....	361
Configuring Microsoft IIS monitoring.....	362
Configuring the agent on Windows systems.....	362
Configuring the agent by using the silent response file.....	363
Changing the user account.....	364
Configuring Microsoft Office 365 monitoring.....	364
Enabling remote execution policy for PowerShell.....	365
Verifying reachability of configured users.....	365
Configuring the agent on Windows systems.....	366
Configuring the agent by using the silent response file.....	367
Changing the user account.....	368
Monitoring the Skype's quality of service.....	368
Configuring local environment variables.....	369
Configuring Microsoft SharePoint Server monitoring.....	370
Changing the user account.....	371
Running Monitoring Agent for Microsoft SharePoint Server by a non-admin user.....	371
Local Security Policy permissions.....	372
Configuring Microsoft SQL Server monitoring.....	373
Creating a user and granting permissions.....	373
Running as a non-administrator user.....	378
Selecting the databases for monitoring.....	380
Configuring local environment variables.....	389
Configuring the agent by using the silent response file.....	395
Running the agent in a cluster environment.....	396
Configuring the agent by using the cluster utility.....	398
Configuring multiple collations for ERRORLOG file.....	399
Configuring historical job count for monitoring.....	400
Configuring MongoDB monitoring.....	401
Configuring the agent with default settings.....	403
Configuring the agent by responding to prompts.....	403
Configuring the agent by using the silent response file.....	404
Configuring MySQL monitoring.....	406
Configuring the agent on Windows systems.....	406
Configuring the agent on Linux systems.....	407
Configuring the agent by using the silent response file.....	407
Configuring NetApp Storage monitoring.....	408
Downloading and installing the NetApp Manageability SDK JAR file.....	409
Configuring the agent on Windows systems.....	410
Configuring the agent by using the silent response file.....	410
Configuring the agent on Linux systems.....	411
Configuration parameters for the data provider.....	412
Configuration parameters for the OnCommand Unified Manager.....	413
Configuration parameters for the OnCommand API Service.....	414
Configuring Oracle Database monitoring.....	414
Configuring the agent on Windows systems.....	416
Configuring the agent by responding to prompts.....	419
Configuring the agent by using the silent response file.....	423
Granting privileges to the Oracle Database agent user.....	425
Configuring OS monitoring.....	428
Run OS agents as a non-root user.....	428
Configure OS agent log file monitoring.....	430
Configuring Linux OS Agent file system data collection.....	450
Configuring PostgreSQL monitoring.....	450

Configuring the agent on Windows systems.....	451
Configuring the agent on Linux systems.....	452
Configuring the agent by using the silent response file.....	453
Configuring RabbitMQ monitoring .....	454
Configuring the agent on Windows systems.....	454
Configuring the agent on Linux systems.....	454
Configuring the agent by using the silent response file.....	455
Configuration parameters for the agent.....	456
Configuring SAP monitoring.....	456
Configuring the agent on Windows systems.....	457
Configuring the agent on Linux or AIX systems.....	458
Configuring the agent by using the silent response file.....	459
Configuration parameters of the agent.....	459
Importing the ABAP transport on the SAP system.....	462
Verifying agent configuration.....	469
Advanced installation and configuration of the SAP agent.....	473
Configuring SAP HANA Database monitoring.....	485
Configuring the agent on Windows systems.....	486
Configuring the agent on Linux and AIX systems.....	487
Configuring the agent by using the silent response file.....	487
Configuring SAP NetWeaver Java Stack monitoring.....	488
Configuring the agent on Windows systems.....	489
Configuring the agent on Linux or AIX systems.....	489
Configuring the agent by using the silent response file.....	490
Configuration parameters of the agent.....	491
Configuring Skype for Business Server monitoring.....	491
Permissions and access rights for a non-administrator user.....	492
Configuring the agent on Windows systems.....	493
Configuring the agent by using the silent response file.....	494
Changing the user account.....	494
Configuration parameters for the agent.....	495
Configuring Sterling Connect Direct monitoring.....	496
Configuring the agent on Windows systems.....	497
Configuring the agent on Linux systems.....	497
Configuring the agent by using the silent response file.....	498
Configuration parameters of the agent.....	499
Configuring Sterling File Gateway monitoring.....	499
Installing the B2B REST API.....	500
Configuring the Sterling File Gateway agent on Windows systems.....	500
Configuring the Sterling File Gateway agent on Linux systems.....	501
Configuring Sterling File Gateway agent by using the silent response file.....	501
Configuring agent environment variables for the data provider on Linux.....	502
Configuring agent environment variables for the data provider on Windows.....	503
Environment variables for the data provider.....	503
Configuration parameters for the B2B API details.....	504
Configuration parameters for database details.....	505
Configuration parameters for the Java API.....	505
Configuring Sybase Server monitoring.....	505
Granting permissions.....	506
Configuring the agent by using the command line interface.....	507
Configuring the agent by using the silent response file.....	509
Disabling dirty reads for query.....	511
Configuring Tomcat Monitoring .....	511
Configuring Tomcat agent with the default settings.....	512
Configuring Tomcat agent on Windows systems.....	512
Configuring Tomcat agent on Linux systems.....	513
Configuring Tomcat agent using silent response file.....	514
Configuring VMware VI monitoring.....	515



Sizing and planning the VMware VI agent deployment.....	516
Enabling SSL communication with VMware VI data sources.....	516
Configuring the agent on Windows systems.....	518
Configuring the agent by responding to prompts.....	518
Configuring the agent by using the silent response file.....	519
Configuration parameters for the data source.....	520
Configuration parameters for the data provider.....	521
Increasing the Java heap size.....	522
Configuring WebSphere Applications monitoring.....	523
Fast track: Configuring the data collector for WebSphere Applications agent.....	523
Checking user access requirements.....	525
Configuring the data collector with the simple configuration utility.....	526
Configuring or reconfiguring the data collector with full configuration utilities.....	527
Manually configure the data collector if the configuration utilities fail.....	542
Restoring the application server configuration from a backup.....	546
Configuring WebSphere Infrastructure Manager monitoring.....	547
Configuring WebSphere MQ monitoring.....	548
Authorizing the user IDs to run the agent.....	548
Configuring IBM MQ (WebSphere MQ) for data enablement.....	550
Configuring the IBM MQ(formerly WebSphere MQ) agent.....	551
<b>Chapter 15. Deploying ICAM Data Collectors.....</b>	<b>555</b>
Kubernetes data collector.....	557
Configuring Kubernetes monitoring.....	557
Configuring Kubernetes monitoring without Ansible.....	562
Configuring Kubernetes monitoring without Helm.....	566
Checking the Kubernetes installation logs.....	569
Kubernetes metrics for thresholds.....	570
Uninstalling the Kubernetes data collector.....	574
Kubernetes data collector issues.....	575
Runtime data collectors.....	577
Common topics.....	577
Configuring Go application monitoring .....	579
Configuring J2SE application monitoring.....	584
Configuring Liberty application monitoring.....	590
Configuring Node.js application monitoring.....	601
Configuring Python application monitoring .....	606
Configuring Ruby application monitoring .....	617
Synthetics PoP.....	620
Installing Synthetics PoP.....	620
Creating a REST API test.....	623
Creating a Scripting REST API synthetic test.....	627
Creating a web page synthetic test.....	632
Creating a Selenium script test.....	633
Event generation.....	639
Viewing Synthetic test results.....	641
Deleting the Synthetics PoP.....	642
Monitoring Agent for HMC.....	643
Installing and configuring the HMC agent.....	643
Uninstalling the HMC agent.....	644
<b>Chapter 16. Deploying Unified Agent.....</b>	<b>645</b>
Overview of Unified Agent.....	645
Preparing the deployment of Unified Agent.....	646
Enabling the NGINX monitoring interface.....	648
Authorizing the plug-ins to access Kubernetes resources.....	651
Installing and configuring the Unified Agent.....	652

Configuring Jaeger and Zipkin in Unified Agent.....	654
Configuring NGINX monitoring in Unified Agent .....	655
Configuring Redis monitoring in Unified Agent .....	656
Configuring IBM API Connect monitoring in Unified Agent .....	656
Configuring IBM App Connect Enterprise monitoring in Unified Agent .....	656
Configuring DEM in Unified Agent .....	657
Configuring IBM MQ monitoring in Unified Agent .....	657
Configuring OpenShift monitoring in Unified Agent .....	659
Updating the Unified Agent configuration.....	659
Uninstalling the Unified Agent.....	662
<b>Chapter 17. Deploying digital experience monitoring(DEM).....</b>	<b>665</b>
Overview of DEM.....	665
DEM for Liberty applications .....	666
DEM for HTTP server.....	667
Installing the UA plug-in for DEM.....	668
Installing and configuring the DEM plug-in for HTTP Server.....	670
Disabling DEM on HTTP server.....	673
<b>Chapter 18. Integrating with other products.....</b>	<b>675</b>
Configuring monitoring data sources.....	675
Integrating with IBM Tivoli Monitoring agents.....	675
Integrating with Cloud APM, Private agents.....	688
Configuring incoming event sources.....	692
Creating custom event sources with JSON.....	692
Configuring Amazon Web Services (AWS) as an event source.....	693
Configuring AppDynamics as an event source.....	694
Configuring Datadog as an event source.....	696
Configuring Dynatrace as an event source.....	697
Configuring Elasticsearch as an event source.....	698
Configuring Jenkins as an event source.....	700
Configuring Logstash as an event source.....	701
Configuring Microsoft Azure as an event source.....	703
Configuring Nagios XI as an event source.....	703
Configuring Netcool/OMNIbus as an event source.....	705
Configuring New Relic as an event source.....	705
Configuring Pingdom as an event source.....	707
Configuring Prometheus as an event source.....	708
Configuring SolarWinds Orion as an event source.....	712
Configuring Splunk Enterprise as an event source.....	714
Configuring IBM UrbanCode Deploy as an event source.....	716
Configuring outgoing event destinations.....	717
Sending incident details to Alert Notification.....	717
Configuring Netcool/OMNIbus as an outgoing incident source.....	718
Sending incident details to Slack channels.....	720
Sending incident details using outgoing webhooks.....	721
Sending incident details to Microsoft Teams.....	722
Sending incident details to ServiceNow.....	723
Sending incident details to GitHub.....	724
Sending incident details to Stride.....	725
Sending incident details to Watson Workspace.....	726
<b>Chapter 19. Administering.....</b>	<b>727</b>
Setting up event policies.....	727
Events and incidents.....	728
Example: Changing event information through enrichment.....	731
Example: Adding to event information through enrichment.....	732

Example: Correlating events through enrichment.....	733
Creating lookup tables.....	734
Example: Enriching event information using lookup tables.....	735
Enriching event information by adding custom attributes .....	737
Example: Suppressing temporary high memory usage warnings.....	738
Example: Preventing certain events from forming incidents.....	739
Example: Setting runbook for disk full events .....	740
Example: Detecting flapping events.....	741
Managing incident policies.....	742
Example: Assigning top priority incidents automatically.....	743
Example: Escalating incidents automatically.....	744
Example: Setting incident priority.....	745
Managing incidents.....	747
Understanding your incidents at a glance.....	749
Understanding the mean time to resolve incidents.....	752
Starting to work with incidents.....	753
Resolving incidents with runbooks.....	754
Managing thresholds.....	755
Managing resource groups.....	758
Setting up users and groups.....	759
Logging in to the Cloud App Management UI to authenticate as a user.....	759
Activating notifications via SMS and voice.....	760
Defining users and groups.....	761
Setting up groups.....	762
Roles.....	763
Using the REST API.....	763
Managing runbooks.....	764
Data retention and summarization.....	765
Configuring summarization.....	766
<b>Chapter 20. Monitoring.....</b>	<b>769</b>
Viewing your managed resources.....	769
Resource dashboard.....	770
Service dependencies topology view.....	775
Transaction tracking.....	780
Monitoring the status of your Tivoli Monitoring data providers.....	782
Viewing Tivoli Monitoring data providers.....	783
ICAM agent and data collector metrics.....	783
Db2 agent metrics.....	784
mySAP agent metrics.....	956
SAP HANA DATABASE agent metrics.....	1117
Sybase Server agent metrics.....	1143
Windows OS agent metrics.....	1204
Kubernetes resources, metrics, and events.....	1373
<b>Chapter 21. Upgrading.....</b>	<b>1387</b>
Upgrading your agents.....	1387
Upgrading your ICAM Agents.....	1387
Upgrading your IBM Tivoli Monitoring agents.....	1388
Upgrading your Cloud APM agents.....	1391
Upgrading your data collectors.....	1392
Upgrading the Synthetics PoP server.....	1393
<b>Chapter 22. Troubleshooting and support.....</b>	<b>1395</b>
Troubleshooting installation and upgrade.....	1395
Installation fails due to permissions error on pods on a Red Hat OpenShift Kubernetes Server.....	1395

PPA installation file fails to load into the IBM Cloud Private repository with Docker authentication error.....	1398
Cloud App Management installation fails with InvalidImageName error.....	1400
Common deployment errors.....	1401
502 Bad Gateway error on IBM Cloud Private .....	1402
IBM Multicloud Manager hub cluster fails to import a cluster with Cloud App Management server installed .....	1403
When scaling up ibm-redis, the ibm-redis-server pods crash.....	1403
Freeing up resources for large Pod scheduling.....	1404
CEM datalayer pod repeatedly tries to restart and fails to start.....	1406
The post-install-setup.sh script runs with the timeout error on Red Hat OpenShift environment.....	1406
Deleting the IBM Cloud Private service instance after the Cloud App Management server is uninstalled.....	1407
Troubleshooting agents.....	1408
Db2 Monitoring.....	1408
Microsoft Active Directory Monitoring.....	1409
PostgreSQL Monitoring.....	1410
Microsoft .NET Monitoring.....	1411
Troubleshooting data collectors.....	1411
Fails to trigger the threshold of Request Rate after upgrading Node.js data collector.....	1411
Troubleshooting the console and dashboards.....	1412
Users are unable to start the Application Monitoring UI. Users can see ERR_TOO_MANY_REDIRECTS in browser.....	1412
IBM Cloud Pak for Multicloud Management email link not working.....	1412
Incidents fail to load in All Incidents tab.....	1413
Resources view cannot display due to internalOAuthError message.....	1414
IBM Tivoli Monitoring (ITM) threshold severities are incorrect in the UI.....	1415
Dashboards for Kubernetes resources are empty after Cloud App Management server upgrade.....	1415
Troubleshooting general issues.....	1415
<b>Events are being routed to the wrong account or team.....</b>	<b>1416</b>
Error message when deleting a service instance.....	1416
Permission issue with Docker Version 18.03 with Ubuntu 16.04 LTS.....	1416
Providing feedback.....	1416
Submitting an IBM Cloud App Management Request for Enhancement (RFE).....	1416
Support.....	1417
Collecting the server logs for IBM Support.....	1417
Collecting monitoring agent logs for IBM Support.....	1418
Contacting IBM Support.....	1418
<b>Chapter 23. Notices.....</b>	<b>1421</b>
Trademarks.....	1422
Terms and conditions for product documentation.....	1423
IBM Online Privacy Statement.....	1423
<b>Accessibility features.....</b>	<b>1431</b>
<b>Notices.....</b>	<b>1433</b>
Trademarks.....	1434
Terms and conditions for product documentation.....	1435
IBM Online Privacy Statement.....	1435

---

# Chapter 1. What's new

New features, capabilities, and coverage are available in the latest release.

For information about the agent version in each release or refresh, see [“Change history”](#) on page 52.

For releases after version 2019.4.0.2, IBM Cloud App Management will not have standalone releases. It will be released as the Monitoring module in IBM Cloud Pak for Multicloud Management. For what's new in future releases, go to [IBM Cloud Pak for Multicloud Management Knowledge Center](#).

## What's new for Version 2019.4.0.2

### 2019.4.0.2

#### New agents

##### CouchDB agent

The Monitoring Agent for CouchDB offers a central point of management for your CouchDB environment or application. The software provides a comprehensive means for gathering the information required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Monitoring Agent for CouchDB you can easily collect and analyze CouchDB specific information. For more information, see [“Configuring CouchDB monitoring”](#) on page 266.

##### Sterling Connect Direct agent

The Monitoring Agent for Sterling Connect Direct monitors the health and performance of Connect Direct nodes in your organization. By using the Sterling Connect Direct agent, you can easily analyze the file transfer activity within and between organizations. For information about configuring the agent after installation, see [“Configuring Sterling Connect Direct monitoring”](#) on page 496.

##### Sterling File Gateway agent

The Monitoring Agent for Sterling File Gateway monitors the Sterling File Gateway application, which is used for transferring files between internal and external partners by using different protocols, different file naming conventions, and different file formats. It also supports the remote monitoring feature. For information about configuring the agent after installation, see [“Configuring Sterling File Gateway monitoring”](#) on page 499.

#### Agent enhancement

##### Microsoft .NET agent

The Microsoft .NET agent is enhanced with the following features:

- Added the following new charts:
  - External Resource Utilization (history)
  - GC Memory Usage (history)
  - GC Collections (history)
  - Thread Queue Length (history)
- Added new Queue Length attribute in the Threads widget, which displays the total number of threads that waited to acquire a managed lock since the application was started.
- Supports Golden Signals that include the following measures: Latency, Saturation, Error, and Traffic.

##### Microsoft Active Directory agent

The Microsoft Active Directory agent is enhanced with the following features:

- Support Golden Signals that include the following measures: Latency, Saturation, Error, and Traffic.

#### **Microsoft SQL Server agent**

The Microsoft SQL Server agent is enhanced with the following features:

- Supports Golden Signals that include the following measures: Latency, Saturation, Error, and Traffic.

#### **Microsoft Hyper-V Server agent**

The Microsoft Hyper-V Server agent is enhanced with the following features:

- Added the Processor Load widget that provides details about the percentage of average processor load of the selected virtual machine.
- Added following new attributes in the Virtual\_Machine\_Details widget that provides information about virtual machines that are hosted on the Hyper-V server:
  - HeartBeat
  - EnabledState
  - HealthState
  - VHDFilePath

#### **Microsoft SharePoint Server agent**

The Microsoft SharePoint Server agent is enhanced with the following features:

- Added the Trace log widget that shows trace log events of SharePoint Server.
- Added the following new group widgets that provide details about the uptime of web application and its associated service applications:
  - Web Application Uptime
  - Service Applications Status
  - Trace Log Details

#### **Tomcat agent**

Added new widget named Signals that shows the golden signals such as latency, traffic and error for application requests.

### **What's new for Version 2019.4.0.1**

#### **IBM Cloud App Management**

##### **Transaction Tracking is enabled for WebSphere® Applications agent**

Enabled Transaction Tracking for the WebSphere Applications agent. For more information, see [“Transaction tracking”](#) on page 780.

**Note:** You need to apply IBM Cloud App Management 2019.4.0-IBM-ICAM-SERVER-IF0002 to view transaction tracking data of WebSphere Applications agent. For readme of this fix, see [2019.4.0-IBM-ICAM-SERVER-IF0002 Readme](#).

##### **Adding Oracle HTTP server support for enabling DEM on HTTP server**

In addition to IBM HTTP Server and Apache HTTP Server, Oracle HTTP server is supported to enable DEM. For more information, see [“DEM for HTTP server”](#) on page 667.

#### **New agents**

##### **Monitoring Agent for Cassandra**

The Monitoring Agent for Cassandra provides you with the capability to monitor the health and performance of Cassandra cluster resources, such as nodes, keyspaces and column families. For information about configuring the agent after installation, see [“Configuring Cassandra monitoring”](#) on page 257.

## Agent enhancement

### Microsoft Cluster Server agent

Added new **Cluster Shared Disks** widget that provides details about the shared disks of Cluster Server, such as disk name and disk state.

### Monitoring Agent for SAP Applications

Added new widget named **Latency** that provides details about latency experienced with respect to CPU time, response time, database time and GUI net time.

### Monitoring Agent for Skype for Business Server

Added the following new widgets:

- Top 5 most active users
- Server failure count
- Failed user count
- Quality of experience
- Usage summary

## What's new for Version 2019.4.0

### IBM Cloud App Management

#### Transaction Tracking

Enabled Transaction Tracking to provide an ability to troubleshoot an issue by looking at a specific request flow through the system and quickly identifying the bottleneck which caused the bad experience for that client. The transaction tracking feature enables topology views and instance level transaction monitoring. By distributed tracking infrastructure, transaction tracking can detect bottleneck issues including latency problems and errors, and filter or sort traces based on application. Transaction tracking can also filter views based on length of trace, timestamp, interactions, errors and transaction comparisons. For more information, see [“Transaction tracking”](#) on page 780.

#### Entitled Registry Installation

Entitled registry is an efficient installation experience using industry best docker-based processes for image management and optimization. Entitled software, IBM Cloud App Management in this case is stored in an IBM Cloud Container Registry `cp.icr.io` domain. To access the IBM Cloud App Management installation image in this domain, you must create an image pull secret with an entitlement key for your cluster and add this image pull secret to the Kubernetes service account of each namespace where you want to deploy the IBM Cloud App Management entitled software. After this setup, you can select the Cloud App Management Helm chart from the **IBM-entitled-charts** in the catalog where it can be quickly installed.

For more information about IBM Cloud App Management with IBM Cloud Pak for Multicloud Management online installations using Entitled registry, see [“Online installation of IBM Cloud App Management with IBM Cloud Pak for Multicloud Management”](#) on page 96 topics.

#### Incident resolution

Navigation enhancements were added to facilitate incident resolution: While investigating an incident, you can select the Event resource page from the Events tab to see a timeline of metrics in context of when the problem surfaced. As well, transaction traces reflect the API that drove the incident.

#### OpenShift monitoring

You can monitor OpenShift route traffic performance and router performance by deploying the Unified Agent plug-in for OpenShift. The plug-in monitors each route response time, volume and error, and also integrates with Kubernetes data collector to enable exploring the associated services and application data. For more information, see [“Configuring OpenShift monitoring in Unified Agent ”](#) on page 659.

### **Adding HTTP server support for digital experience monitoring (DEM)**

Besides Liberty applications, you can now enable DEM on HTTP server by installing the DEM plug-in for HTTP Server. The DEM plug-in for HTTP Server passively collects data on how actual users are interacting with and experiencing your application. This is achieved through instrumenting the application or injecting code on the page to collect metrics. With DEM and Transaction Tracking, the DEM plug-in for HTTP Server can monitor web application performance from the browser to the line of code. For more information, see [Chapter 17, “Deploying digital experience monitoring\(DEM\),”](#) on page 665.

### **Adding geolocation information for digital experience monitoring (DEM)**

After applying 2019.4.0-IBM-ICAM-SERVER-IF0001, you can view geolocation information of the real users. In the **Browser** dashboard, click **Filter** to see the country, region, and city information of each application real user.

**Note:** To enable this feature, you must do the following steps:

- Apply 2019.4.0-IBM-ICAM-SERVER-IF0001. For details, see [IBM Cloud App Management 2019.4.0 2019.4.0-IBM-ICAM-SERVER-IF0001 Readme](#).
- Meet the per-requisites as stated in [DEM overview](#).

## **New agents and data collectors**

### **Citrix VDI agent**

The Monitoring Agent for Citrix Virtual Desktop Infrastructure monitors the following functions: Citrix XenDesktop component, Event log and alerts, and Citrix XenDesktop services. Additionally, you can view the Load Index Summary metrics performance data for Citrix XenApp and XenDesktop. You can diagnose problematic login times by viewing the performance data for the login steps. For more information, see [“Configuring Citrix Virtual Desktop Infrastructure monitoring”](#) on page 259.

### **Go data collector**

The Go data collector can provide you with visibility and control of your Go applications, and help you ensure optimal performance and efficient use of resources. You can reduce and prevent application crashes and slowdowns around the clock, as the data collector assists you in detecting, diagnosing and isolating performance issues. For more information, see [“Configuring Go application monitoring ”](#) on page 579.

### **HMC agent**

The Monitoring Agent for HMC provides you with the capability to monitor the Hardware Management Console (HMC). The HMC agent monitors the availability and health of HMC resources such as CPU, memory, storage, and network. It collects the following metrics: HMC, Managed Server(CEC), LPAR, VIOS, CPUPool, VSCSI, FibreChannel, and NPIV and sends these metrics to the Cloud App Management server. The supported HMC versions are V8.2 to V8.7, and V9.1. For information about configuring the agent after installation, see [Installing and Configuring the HMC agent](#).

### **Microsoft Active Directory agent**

The Monitoring Agent for Microsoft Active Directory provides capabilities to monitor the Active Directory in your organization. You can use the agent to collect and analyze information that is specific to Active Directory. For information about configuring the agent after installation, see [“Configuring Microsoft Active Directory monitoring”](#) on page 340.

### **RabbitMQ agent**

The Monitoring Agent for RabbitMQ provides you with the capability to monitor the RabbitMQ cluster. You can collect and analyze information about the nodes, queues, and channels of the RabbitMQ cluster. For information about configuring the agent after installation, see [“Configuring RabbitMQ monitoring ”](#) on page 454.

### **Sybase agent**

The Sybase agent offers a central point of management for distributed databases. It collects the required information for database and system administrators to examine the performance of the



Sybase server system, detect problems early and prevent them. For information about configuring the agent after installation, see [“Configuring Sybase Server monitoring”](#) on page 505.

### **Ruby data collector**

The Ruby data collector can provide you with visibility and control of the Ruby application, and help you ensure optimal performance and efficient use of resources. For more information, see [“Configuring Ruby application monitoring”](#) on page 617.

### **WebLogic agent**

The Monitoring Agent for WebLogic provides you with a central point of monitoring for the health, availability, and performance of your WebLogic server environment. The agent displays a comprehensive set of metrics to help you make informed decisions about your WebLogic resources, including Java™ virtual machines (JVMs), Java messaging service (JMS), Java Database Connectivity (JDBC). For information about configuring the agent after installation, see [“Configuring WebLogic monitoring”](#) on page 270.

## **Agent and data collector enhancements**

### **Expanded platform support for agents**

- Solaris x86-64

### **Db2® agent**

- Added support for Solaris x86-64.
- Added support for RHEL 8 on x86-64 and Power Linux Little Endian (pLinux LE) (64 bit).

### **Hadoop agent**

- Added new widgets to monitor Job Details parameters: Failed Map and Reduce tasks, and Map Output Record spills
- Added legends for the charts having more than one metric.

### **Microsoft Exchange Server agent**

- Added support for Microsoft Exchange Server 2013, 2016 and 2019.
- Added resource type `Microsoft Exchange Server` for Microsoft Exchange Server 2007 and 2010.
- Added resource type `Microsoft Exchange Server 2k13` for Microsoft Exchange Server 2013, 2016 and 2019.

### **Python data collector**

Added the support of uWSGI and mod\_wsgi/Apache httpd. For more information, see [“Configuring Python application monitoring”](#) on page 606.

### **SAP agent**

A line chart is added on the UI depicting the count of errors that occur on the SAP system. Error counts shown in the line chart are:

- Kernel Errors
- Database Interface
- ABAP Programming Errors
- Installation Errors
- Resource Shortage Errors

### **SAP HANA Database agent**

A line chart is added on the UI depicting the count of errors that occur on the SAP HANA Database. Error counts shown in the line chart are:

- Network error
- CPU Errors
- Storage Errors

- Administrator / Configuration Errors
- Memory errors

## What's new for Version 2019.3.0.1

### New agents

#### MariaDB agent

The Monitoring Agent for MariaDB offers a central point of management for your MariaDB environment or application. The software provides a comprehensive means for gathering the information required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Monitoring Agent for MariaDB you can easily collect and analyze MariaDB specific information. For information about configuring the agent after installation, see [“Configuring MariaDB monitoring” on page 335](#)

#### Amazon EC2 agent

The Monitoring Agent for Amazon EC2 offers a central point of management for your Amazon EC2 environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Amazon EC2 agent you can easily collect and analyze Amazon EC2 specific information. For more information, see [“Configuring Amazon EC2 monitoring” on page 232](#).

#### Amazon ELB agent

The Monitoring Agent for AWS Elastic Load Balancer offers a central point of management for your AWS Elastic Load Balancer environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Amazon ELB agent you can easily collect and analyze AWS Elastic Load Balancer specific information. For more information, see [“Configuring AWS Elastic Load Balancer monitoring” on page 238](#).

#### Azure Compute agent

The Monitoring Agent for Azure Compute offers a central point of management for your Azure Compute environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Azure Compute agent you can easily collect and analyze Azure Compute specific information.. For more information, see [“Configuring Azure Compute monitoring” on page 242](#).

### Agent enhancement

#### MySQL agent

The MySQL agent is enhanced with the following features:

- Displayed more descriptive help text or tooltip for the following widgets:
  - Active Connections Percentage Widget
  - Database Size Information Widget
- Expanded the column of Processlist Details Widget for better visibility and readability.
- Agent is able to collect data consistently after server restarts.

## What's new for Version 2019.3.0

### New agents

#### Cisco UCS agent

The Monitoring Agent for Cisco UCS provides you with an environment to monitor the health, network, and performance of Cisco Unified Computing Systems (UCS). The Cisco UCS agent provides a comprehensive way for collecting and analyzing information that is specific to Cisco UCS and required to detect problems early and prevent them.

For information about configuring the agent after installation, see [“Configuring Cisco Unified Computing System \(UCS\) monitoring” on page 250.](#)

#### Microsoft Cluster Server agent

The Monitoring Agent for Microsoft Cluster Server provides capabilities to monitor the Microsoft Cluster Server in your organization. You can use the Microsoft Cluster Server agent to collect information that is related to cluster resource availability, such as cluster level, cluster nodes, cluster resource groups, cluster resources, and cluster networks. The agent also provides statistics for cluster resources usage, such as processor usage, memory usage, disk usage, and network usage.

For information about configuring the agent after installation, see [“Configuring Microsoft Cluster Server monitoring” on page 344](#)

#### Microsoft SharePoint Server agent

The Monitoring Agent for Microsoft SharePoint Server provides you with the environment to monitor the availability, events, and performance of the Microsoft SharePoint Server. Use this agent to gather data from the Microsoft SharePoint Server and manage operations.

For information about configuring the agent after installation, see [“Configuring Microsoft SharePoint Server monitoring” on page 370.](#)

#### MongoDB agent

The Monitoring Agent for MongoDB provides monitoring capabilities for the usage, status, and performance of the MongoDB deployment. You can collect and analyze information such as database capacity usage, percentage of connections open, memory usage, instance status, and response time in visualized dashboards.

For information about configuring the agent after installation, see [“Configuring MongoDB monitoring” on page 401.](#)

#### MySQL agent

The Monitoring Agent for MySQL provides monitoring capabilities for the status, usage, and performance of the MySQL deployment based on the 2 top level resources classification.

- MySQL Database

On the Database Resource, you can view and analyze the information specific to different databases of your MySQL Server, such as Table Count, Database Size, ProcessList Details, Events data and others.

- MySQL Instance

On the Resource page, you can view all the instances for MySQL and analyze information such as Percentage of Active Connections, Slow Queries, Bytes Received vs Sent, Error Details, CPU, Memory Usage and others.

For information about configuring the agent after installation, see [“Configuring MySQL monitoring” on page 406.](#)

#### PostgreSQL agent

The Monitoring Agent for PostgreSQL monitors the PostgreSQL database by collecting PostgreSQL metrics through a JDBC driver. The agent provides data about system resource usage, database capacity, connections that are used, individual status of running instances, statistics for operations, response time for SQL query statements, database size details, and lock information.

For information about configuring the agent after installation, see [“Configuring PostgreSQL monitoring”](#) on page 450.

### **SAP NetWeaver Java Stack agent**

The Monitoring Agent for SAP NetWeaver Java Stack monitors the availability, resource usage, and performance of the SAP NetWeaver Java Stack. The agent can monitor SAP NetWeaver Java Stack deployment scenarios such as single host - single instance, single host - multiple instances, multiple hosts - single instances, and multiple hosts - multiple instances. You can analyze the information that the agent collects and take appropriate actions to resolve issues in the SAP NetWeaver Java Stack.

For information about configuring the agent after installation, see [“Configuring SAP NetWeaver Java Stack monitoring”](#) on page 488.

## **Agent and data collector enhancements**

### **Agents and data collectors can be deployed for IBM Multicloud Manager**

You can install and configure agents and data collectors after you deploy the Cloud App Management Klusterlet for IBM Multicloud Manager. For more information, see [“Deploying agents and data collectors for IBM Cloud Pak for Multicloud Management”](#) on page 140.

### **Expanded platform support for agents**

- Solaris Sparc
- Linux for Power® Systems (pLinux)
- Linux for Power Systems Little Endian (pLinux LE)

### **Db2 agent**

- Added support for Solaris Sparc 11.
- Added support for RHEL 7 on Power Linux Little Endian (pLinux LE) (64 bit).
- Added support for Db2 Server Version 11.5.

### **Unified Agent**

More plug-ins have been enabled on the Unified Agent. By deploying the Unified Agent, you can monitor IBM App Connect Enterprise and IBM MQ that are deployed in IBM Cloud Private environment, NGINX and Redis workloads, and IBM API Connect application. For more information, see [Chapter 16, “Deploying Unified Agent,”](#) on page 645.

### **Kubernetes data collector**

In previous releases, you deployed the cloud data collector for monitoring the applications in your Kubernetes environment, including NGINX and Redis. The Unified Agent is now used to deploy NGINX and Redis plug-ins. The cloud data collector has been renamed to Kubernetes data collector and you no longer configure NGINX or Redis monitoring with the Kubernetes data collector.

### **Runtime data collectors enhancements**

- Added the support of Flask framework in Python data collector.
- Added OpenTracing sampling and Latency sampling in runtime data collectors including Node.js data collector, Liberty data collector, J2SE data collector, and Python data collector.

### **Digital experience monitoring (DEM, previously called RUM)**

RUM (Real User Monitoring) is renamed as DEM and has the following enhancements:

- A more simple deployment method of enabling and disabling DEM. For more information, see [“DEM for Liberty applications ”](#) on page 597.
- Added the **Browser** dashboard that can be navigated from the **Service dependencies** topology or the **Related resources** widget.

- In the **Browser** dashboard, you can view real user experience data, including golden signal and latency breaking down.
- Browser view is context-sensitive. You can view browser metrics of the Kubernetes service where you navigate from.
- In the **Browser** dashboard, you can filter by browser type.

### New Synthetic test types

The following three new synthetic test types are now available:

Scripting REST API synthetic test: Test and monitor a number of REST APIs in a sequence using a `node.js` script.

Web page synthetic test: Test a single web page for availability and browser response time.

Selenium script synthetic test: Test simulated user interactions with your web application.

## Cloud App Management server

### Integration with IBM Cloud Pak for Multicloud Management

The IBM Cloud App Management integration with the IBM Cloud Pak for Multicloud Management provides you with more application and cluster visibility across the enterprise to any public or private cloud. You can improve your IT and application operations management with increased flexibility and cost savings, and intelligent data analysis driven by predictive signals.

### Support for IBM Power LE with IBM Cloud Pak for Multicloud Management

The Cloud App Management server now supports running on IBM Power LE as part of the IBM Cloud Pak for Multicloud Management.

### Data retention

Data retention refers to the number of days that data samples are saved for viewing in the **Resources** dashboards. In earlier releases the default setting was 32 days. Now the default setting is 8 days. You can configure a different data retention setting during Cloud App Management server installation or upgrade. For more information, see [“About data retention and summarization” on page 765](#).

### Data summarization

Summarizing data enables you to perform historical analysis of data over time, examine trends, and do high level capacity planning. During Cloud App Management server installation or upgrade, you can now enable summarization for agents that support it: Linux KVM agent, Linux OS agent, UNIX OS agent, and VMware VI agent. For more information, see [“About data retention and summarization” on page 765](#).

## Cloud App Management console

### Resource groups

You can now create resource groups based on resource type, such as Kubernetes Service, or by selecting individual managed resources and then assign to thresholds. For more information, see [“Managing resource groups” on page 758](#).

### Thresholds

The **missing**, **match**, and **not match** relational operators were added for threshold conditions that use text metrics. The **average** and **count** functions are available for aggregate metrics. You can also define a reflex action to take place after an event is opened. For more information, see [“Managing thresholds” on page 755](#).

### Resources view

The Cloud Resources tab has been removed. The ICAM Data Collectors dashboards are now accessed through the Resources tab.

## Topology view for Kubernetes services with dependencies

In the Resource Dashboard, for Kubernetes services with dependencies, you can now navigate from the Service Dependencies widget to a Service Dependencies topology view. For more information, see [“Service dependencies topology view”](#) on page 775.

## Monitoring IBM Tivoli® Monitoring data providers

If you have issues with your Tivoli Monitoring resources producing data, there might be an issue with the data providers for these resources. You can verify the status (online or offline) of the resource data providers from the **Resource** dashboard.

You can also quickly view other information about Tivoli Monitoring data providers in one view on the **Monitoring Data Providers** page in the Cloud App Management console. For more information, see [“Monitoring the status of your Tivoli Monitoring data providers”](#) on page 782.

## What's new for Version 2019.2.1.2

### Unified Agent

The Unified Agent is an agent for collecting, processing, aggregating, and writing metrics to your IBM Cloud App Management environment. It is based on Telegraf, and supports receiving open tracing workloads including Jaeger and Zipkin. For more information, see [Chapter 16, “Deploying Unified Agent,”](#) on page 645.

### New agents

#### Microsoft Exchange Server agent

The Monitoring Agent for Microsoft Exchange Server provides capabilities to monitor the health, availability, and performance of the Exchange Servers in your organization. You can use the Microsoft Exchange Server agent to collect server-specific information, such as mail traffic, state of mailbox databases and activities of clients. Additionally, the agent provides statistics of cache usage, mail usage, database usage and client activities to help you analyze the performance of Exchange Servers.

For information about configuring the agent after installation, see [“Configuring Microsoft Exchange Server monitoring”](#) on page 346.

#### Microsoft Office 365 agent

The Monitoring Agent for Microsoft Office 365 provides capabilities to monitor your Microsoft Office 365 environment or application. You can use the Microsoft Office 365 agent to monitor the health and performance of Office 365 resources, such as the Office 365 subscribed services, Office 365 portal, mailbox users, SharePoint sites, and OneDrive storage.

For information about configuring the agent after installation, see [“Configuring Microsoft Office 365 monitoring”](#) on page 364.

#### NetApp Storage agent

The Monitoring Agent for NetApp Storage provides capabilities to monitor your NetApp storage systems by using the NetApp OnCommand Unified Manager (OCUM). You can use the NetApp Storage agent to monitor the health and performance of ONTAP cluster with event-driven responses and precise representation of historical trends in the Cloud App Management console.

For information about configuring the agent after installation, see [“Configuring NetApp Storage monitoring”](#) on page 408.

## What's new for Version 2019.2.1.1

### Open Liberty support

Added the support of Open Liberty.

### New agents

#### DataStage® agent

The Monitoring Agent for InfoSphere® DataStage offers a central point of management for InfoSphere DataStage application's Service Tier as well as Engine Tier. You can use the InfoSphere DataStage agent to monitor details, such as Job Runs, CPU and Memory of engines, historical trend, status of services, and so on. Information is standardized across the system. You

can monitor multiple engines from a single point. By using the InfoSphere DataStage Application agent you can easily collect and analyze InfoSphere DataStage Application specific information. For information about configuring the agent after installation, see [“Configuring InfoSphere DataStage monitoring”](#) on page 313.

### **Hadoop agent**

The Monitoring Agent for Hadoop provides capabilities to monitor the Hadoop cluster in your organization. You can use the agent to collect and analyze information about the Hadoop cluster, such as status of data nodes and Java™ virtual machine, memory heap and non-heap information, and information about Hadoop nodes, file systems, and queues.

For information about configuring the agent after installation, see [“Configuring Hadoop monitoring”](#) on page 292.

### **Skype for Business Server agent**

The Monitoring Agent for Skype for Business Server provides you with the capability to monitor the Skype for Business Server. You can use the agent to monitor the availability, performance, error log, event log, and historical data of the Business Server.

For information about configuring the agent after installation, see [“Configuring Skype for Business Server monitoring”](#) on page 491.

## **Agent and data collector enhancements**

### **Microsoft IIS agent**

The following new group widgets are added to show Worker Process Details, .Net Memory Management, Network and Connection Statistics, and System-Main Memory Statistics, which helps the administrator to identify problems easily:

- System - Main Memory Statistics
- Total Method Requests per second
- Network Statistics
- Connection Statistics
- .Net Memory Management
- Worker Process Details

### **Microsoft SQL Server agent**

The widget Job details is enhanced to display the **Success count** and **Non-success count** based on the configuration of **Maximum job history rows per job**.

### **Python data collector**

Open tracing is enabled for Python data collector by default. You can disable open tracing by following the instructions at [“Customizing the Python data collector”](#) on page 613.

### **Tomcat agent**

A new resource named JVM Runtime is added to enhance the JVM monitoring.

## **What's new for Version 2019.2.1**

### **Red Hat OpenShift support**

You can deploy Cloud App Management running on IBM Cloud Private with Red Hat OpenShift. For more information, see [“Offline: Installing IBM Cloud App Management stand-alone on Red Hat OpenShift”](#) on page 141.

### **IBM certified container**

Cloud App Management has achieved IBM certified container status. This ensures that Cloud App Management meets the enhanced enterprise-grade criteria for security, integration, and workload availability. For more information, see the [Identifying IBM certified containers](#) topic in the IBM Cloud Private Knowledge Center.

## **IBM Multicloud Manager**

Cloud App Management can be integrated with IBM Multicloud Manager. IBM Multicloud Manager allows you to effectively manage multiple cloud environments (public or private) as if they were a single environment. For more information, see [“Installing IBM Cloud App Management with IBM Cloud Pak for Multicloud Management”](#) on page 94.

## **High availability**

Cloud App Management can now be deployed in a high availability environment. For more information, see [“Planning for a high availability installation”](#) on page 148.


## **Backup and restore**

Backup and restore your workload. For more information, see [“Backing up and restoring”](#) on page 177.


## **Real User Monitoring (RUM)**

Cloud App Management has added the support of RUM that collects data on how actual users are interacting with and experiencing web applications. It is achieved through instrumenting the application and injecting code on the page to collect metrics. You can enable RUM for Liberty data collector. For more information, see [“DEM for Liberty applications”](#) on page 597.


## **Event enrichment using lookup tables**

You can use lookup tables to enrich events by correlating attributes in the events with corresponding attributes in the lookup table. Event policies can contain multiple lookup tables. For more information, see the [Creating lookup tables](#) and [Example: Enriching event information using lookup tables](#)  topics in the Cloud Event Management Knowledge Center.


## **Enhanced event forwarding to Netcool/OMNIbus**

Forward events from Cloud App Management to Netcool/OMNIbus with the IBM Secure Gateway and enhanced event policies. For more information, see the [Sending events to Netcool®/OMNIbus via the IBM Secure Gateway](#) and [Setting up event policies](#)  topics in the Cloud Event Management Knowledge Center.


## **Integrate with VMware vSphere (Cloud App Management advanced only)**

Set up VMware vSphere as an event source in Cloud App Management advanced, and start receiving notifications created by VMware vSphere. For more information, see the [Configuring VMware vCenter Server as an event source](#)  topic in the Cloud Event Management Knowledge Center.

## **Enhanced integration with Microsoft Azure (IBM Cloud App Management, Advanced only)**

Azure Log Alert - Log Analytics is now supported in the integration with IBM Cloud App Management, Advanced. For more information, see the [Configuring Microsoft Azure as an event source](#)  topic in the Cloud Event Management Knowledge Center.

## **Runbook Automation triggers**

Define triggers to connect events in Netcool/OMNIbus to runbooks. You can launch a runbook in Cloud App Management from the Web GUI Event Console in Netcool/OMNIbus. For more information, see the [Triggers](#)  topic in the Runbook Automation Knowledge Center.

## **Thresholding enhancements**

You now have the ability to define thresholds on a filtered set of sub-resources and to define complex thresholds on a single resource. You can select a specific sub-resource for a condition, such as a specific disk rather than all disks. You can apply multiple conditions to a threshold with Boolean AND logic or Boolean OR logic (or both).

## **Visualization enhancements**

The Resource and Cloud Resource dashboards present metrics from difference perspectives. The SRE golden signals views were added for release 2019.2.0. The golden signal views are now available for



more resource types. In addition, a **Golden signals** tab is now displayed on the dashboard as well as an **Infrastructure** tab or **Pod network** tab for certain resource types.

## **New agents and data collectors**

### **SAP HANA Database agent**

The Monitoring Agent for SAP HANA Database monitors availability, resource usage, and performance of the SAP HANA database. It can monitor HANA deployment scenarios such as single host - single database, single host - multiple tenant databases, multiple hosts - single database, and multiple hosts - multiple tenant databases. You can analyze the information that the agent collects and take appropriate actions to resolve issues in the SAP HANA Database. For more information about configuring the agent after installation, see [“Configuring SAP HANA Database monitoring”](#) on page 485.

### **WebSphere Infrastructure Manager agent**

The Monitoring Agent for WebSphere Infrastructure Manager monitors the performance of WebSphere Deployment Manager and Node Agent. The WebSphere Infrastructure Manager agent is a multiple instance agent. You must create the first instance and start the agent manually. For more information about configuring the agent after installation, see [“Configuring WebSphere Infrastructure Manager monitoring”](#) on page 547.

### **Python data collector**

The Data Collector for Python monitors your Django based Python applications. Through detecting, diagnosing, and isolating performance issues, the Python data collector helps you ensure optimal performance and efficient use of resources, reduce, and prevent application crashes and slowdowns around the clock. For more information, see [“Configuring Python application monitoring”](#) on page 606.

## **Enhanced agents and data collectors**

### **IBM Integration Bus agent**

Added tolerance support to monitor IBM App Connect Enterprise V11. For more information, see [“Configuring IBM Integration Bus monitoring”](#) on page 304.

### **J2SE data collector**

Added the support of showing metrics for golden signal data including request data.

### **Liberty data collector**

Real User Monitoring(RUM) can be enabled for Liberty data collector to passively collect data about how actual users interact with and experience web applications. For more information, see [“DEM for Liberty applications”](#) on page 597.

For information about the agent version in each release or refresh, see [“Change history”](#) on page 52.

## **Expanded platform support for agents**

### **Red Hat Enterprise Linux (RHEL) 8**

The following ICAM Agents now support RHEL 8. Before installing agents on RHEL 8, be sure to read the [“Specific operating systems”](#) on page 201 section of [“Preinstallation on Linux systems”](#) on page 201.

#### **RHEL 8 on x86-64 (64 bit)**

- Db2 agent
- HTTP Server agent
- Linux KVM agent
- Linux OS agent
- SAP agent
- SAP HANA Database agent

- Tomcat agent
- VMware VI agent
- WebSphere Applications agent
- IBM MQ(formerly WebSphere MQ) agent

#### **RHEL 8 on System z®**

- Db2 agent
- HTTP Server agent
- Linux KVM agent
- Linux OS agent
- WebSphere Applications agent
- IBM MQ(formerly WebSphere MQ) agent

#### **Prerequisite scanner**

The **IGNORE\_PRECHECK\_WARNING** command is now available as an alternative to the **SKIP\_PRECHECK** command. For more information, see [“Bypassing the prerequisite scanner” on page 213.](#)

#### **What's new for Version 2019.2.0.1**

##### **New agents**

##### **Microsoft .NET agent**

The Monitoring Agent for Microsoft .NET offers a central point of management for your Microsoft .NET environment or application. With the Monitoring Agent for Microsoft .NET, you can easily collect and analyze Microsoft .NET specific information from Cloud App Management console. The agent also monitors various applications, services and processes that uses the .Net CLR.

For information about configuring the agent after installation, see [“Configuring Microsoft .NET monitoring” on page 338.](#)

##### **Microsoft IIS agent**

The Monitoring Agent for Microsoft Internet Information Services offers a central point of management for your Microsoft Internet Information Server environment or application. You can use the Microsoft Internet Information Server agent to monitor website details such as request rate, data transfer rate, error statistics, and connections statistics. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Microsoft Internet Information Server agent you can easily collect and analyze Microsoft Internet Information Server specific information.

For information about configuring the agent after installation, see [“Configuring Microsoft IIS monitoring” on page 362.](#)

##### **Microsoft SQL Server agent**

The Monitoring Agent for Microsoft SQL Server offers a central point of monitoring for your Microsoft SQL Server environment or application. You can collect and analyze Microsoft SQL Server specific information, and monitor multiple servers from a single IBM Cloud® App Management console.

For information about configuring the agent after installation, see [“Configuring Microsoft SQL Server monitoring ” on page 373.](#)

##### **SAP agent**

The Monitoring Agent for SAP Applications provides the capability to monitor your SAP system. The SAP agent offers a central point of management for gathering the information to detect problems early, and prevent them. It enables effective systems management across SAP releases, applications, components, and the underlying databases, operating systems, and external interfaces.

For more information about configuring the agent after installation, see [“Configuring SAP monitoring”](#) on page 456.

### **Tomcat agent**

The Monitoring Agent for Tomcat offers a central point of management for your Tomcat environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single IBM Cloud App Management console. By using the Tomcat agent you can easily collect and analyze Tomcat specific information.

For information about configuring the agent after installation, see [“Configuring Tomcat Monitoring”](#) on page 511

## **What's new for Version 2019.2.0**

### **New agents**

#### **Citrix VDI agent**

The Monitoring Agent for Citrix Virtual Desktop Infrastructure monitors the following functions: Citrix XenDesktop component, Event log and alerts, and Citrix XenDesktop services. Additionally, you can view the Load Index Summary metrics performance data for Citrix XenApp and XenDesktop. You can diagnose problematic login times by viewing the performance data for the login steps. For more information, see [“Configuring Citrix Virtual Desktop Infrastructure monitoring”](#) on page 259.

#### **JBoss agent**

The Monitoring Agent for JBoss offers a central point of management for your JBoss environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the JBoss agent you can easily collect and analyze JBoss specific information. For more information, see [“Configuring JBoss monitoring”](#) on page 316.

#### **Microsoft Hyper-V Server agent**

The Monitoring Agent for Microsoft Hyper-V Server provides capability to monitor the availability and performance of all the Hyper-V systems in your organization. The Microsoft Hyper-V Server agent provides configuration information such as the number of virtual machines, the state of the virtual machines, the number of allocated virtual disks, the allocated virtual memory, and so on. Additionally, the agent provides statistics of physical processor usage, memory usage, network usage, logical processor usage, and virtual processor usage.

For information about configuring the agent after installation, see [“Configuring Microsoft Hyper-V monitoring”](#) on page 358.

#### **Monitoring Agent for Linux KVM**

The Monitoring Agent for Linux KVM offers a central point of management for your Linux Kernel-based Virtual Machines environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single IBM Cloud App Management console. By using the Linux KVM agent you can easily collect and analyze Linux Kernel-based Virtual Machines specific information.

For information about configuring the agent after installation, see [“Configuring Linux KVM monitoring”](#) on page 324.

#### **Monitoring Agent for VMware VI**

The Monitoring Agent for VMware VI offers a central point of management for your VMware VI environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single IBM Cloud App Management console. By using the VMware VI agent you can easily collect and analyze VMware specific information.

For information about configuring the agent after installation, see [“Configuring VMware VI monitoring”](#) on page 515.

## **New data collector**

### **J2SE data collector**

The J2SE data collector is a greenfield runtime data collector that monitors the cloud-based Java applications. The J2SE data collector helps you to manage the performance and availability of stand-alone Java applications in IBM Cloud Private.

For information about configuring the data collector, see [“Configuring J2SE application monitoring”](#) on page 584.

## **Data collector enhancements**

### **Kubernetes monitoring**

To further remove incident noise and help you identify root cause, the cloud data collector incident correlation and status propagation have been enhanced.

You can now define custom thresholds for Kubernetes Cluster, Kubernetes Node, and Kubernetes Pod resource types. For more information about thresholds, see [“Managing thresholds”](#) on page 755.

Kubernetes resource types have been added for viewing other facets of your environment. You can select from these new resource types: deployment, job, daemon set, stateful set, application runtime, cron job, and application CRD.

After you select the **View Resources** button for one of the Kubernetes resource types, the list of resources that are displayed provide additional columns. For example, select Kubernetes Service and you see Status, Resource, and Service type, as well as these new columns: Cluster, Namespace, Cluster IP address, and Ports.

The Kubernetes dashboards now present line charts of the SRE four golden signals, as well as single-hop dependencies. For more details, see the [“Resource dashboard”](#) on page 770 information.

### **Synthetics PoP**

You can install and configure multiple Synthetics PoP components, which enable you to create and run synthetic REST API tests. Configure events and alerts in response to slow or unresponsive synthetics test playbacks. In the Cloud App Management console, visualize metrics for your synthetics tests in relation to availability and response time. For more information, see [“Synthetics PoP”](#) on page 620.

## **Documentation enhancement**

A page is created to help you quickly find out the version information and change history for each agent and data collector. See [“Change history”](#) on page 52.

---

## Chapter 2. Known issues and limitations

Review the known issues for IBM Cloud App Management. Additionally, see [Chapter 22, “Troubleshooting and support,”](#) on page 1395 for troubleshooting topics.

- [“Unable to start the Application Monitoring UI because of ERR\\_TOO\\_MANY\\_REDIRECTS error”](#) on page 17
- [“Unable to load root/.rnd into RNG when installing IBM Cloud App Management stand-alone”](#) on page 17

### Unable to start the Application Monitoring UI because of ERR\_TOO\_MANY\_REDIRECTS error

#### Issue

When IBM Cloud App Management 2019.4.0 is installed with IBM Cloud Pak for Multicloud Management, you cannot start the Application Monitoring UI because an `ERR_TOO_MANY_REDIRECTS` error is displayed in your browser. This error is caused by a `302 error` accessing the `https://(hostname)/cemui/launch/auth` URL.

#### Solution

To fix this issue, complete the following steps::

1. Apply the `2019.4.0-IBM-ICAM-SERVER-IF0001` fix. For more information, see the [IBM Cloud App Management 2019.4.0 2019.4.0-IBM-ICAM-SERVER-IF0001](#) README.
2. Start the Application Monitoring UI again.

#### Workaround

It is recommended that you complete the previous solution to completely resolve the issue. However, if you cannot apply the `2019.4.0-IBM-ICAM-SERVER-IF0001` fix currently, the following workaround fixes the issue temporarily and you can start the Application Monitoring UI.

1. Access the **Incidents** page and authenticate. Next, browse to your resources from the **Incidents** page.
2. After the credentials are cached, access the Application Monitoring UI again.

### Unable to load root/.rnd into RNG when installing IBM Cloud App Management stand-alone

#### Issue

Sometimes when you are running the `pre-install.sh` or `make-ca-cert-icam.sh` scripts while you are installing IBM Cloud App Management, you get the following error:

```
Can't load /root/.rnd into RNG
139718552379840:error:2406F079:random number generator:RAND_load_file:Cannot open file:../
crypto/rand/randfile.c:88:Filename=/root/.rnd
Can't load /root/.rnd into RNG
```

#### Solution

To fix this issue, complete the following steps:

1. Install one of the following packages depending on your operating system:
  - Ubuntu: Install `libssl-dev`.
  - Red Hat Enterprise Linux (RHEL): Install `openssl-devel`.
2. Rerun the script that you ran previously when you got the error, which is either `pre-install.sh` or `make-ca-cert-icam.sh`.
3. If the script still fails with the same error again, run the following and then rerun the script again:

```
cd /tmp
rm -rf ibmsecrets
mkdir ibmsecrets
cd ibmsecrets
export RANDFILE=/tmp/ibmsecrets/.rnd
touch $RANDFILE
```

---

## Chapter 3. PDF documentation

A PDF document is available for the topics in this IBM Knowledge Center collection as a User's Guide. References are also available for each agent.

### IBM Knowledge Center

The following PDF document provides Cloud App Management topics from the IBM Knowledge Center in a printable format.

[IBM Cloud App Management User's Guide](#)

## Documentation

You can find information for IBM Cloud App Management in the IBM Knowledge Center.

### IBM Knowledge Center

[IBM Cloud App Management](#) in the IBM Knowledge Center is the official source of technical information for the product.

Information is also available at the following websites:

### [Software Product Compatibility Reports \(SPCR\) tool](#)

You can use the SPCR tool to generate various types of reports that are related to offering and component requirements. Search for IBM Cloud App Management.

## Conventions used in the documentation

Several conventions are used in the documentation for special terms, actions, commands, paths that are dependent on your operating system, and for platform-specific and product-specific information.

### Typeface conventions

The following typeface conventions are used in the documentation:

#### **Bold**

- Lowercase commands, mixed-case commands, parameters, and environment variables that are otherwise difficult to distinguish from the surrounding text
- Interface controls (check boxes, push buttons, radio buttons, spin buttons, fields, folders, icons, list boxes, items inside list boxes, multicolumn lists, containers, menu choices, menu names, tabs, property sheets), labels (such as **Tip:**)
- Keywords and parameters in text

#### *Italic*

- Citations (examples: titles of publications, diskettes, and CDs)
- Words and phrases defined in text (example: a nonswitched line is called a *point-to-point line*)
- Emphasis of words and letters (example: The LUN address must start with the letter *L*.)
- New terms in text, except in a definition list (example: a *view* is a frame in a workspace that contains data.)
- Variables and values you must provide (example: where *myname* represents...)

#### **Monospace**

- Examples and code examples
- File names, directory names, path names, programming keywords, properties, and other elements that are difficult to distinguish from the surrounding text
- Message text and prompts

- Text that you must type
- Values for arguments or command options

### **Bold monospace**

- Command names, and names of macros and utilities that you can type as commands
- Environment variable names in text
- Keywords
- Parameter names in text: API structure parameters, command parameters and arguments, and configuration parameters
- Process names
- Registry variable names in text
- Script names

### **Operating system-dependent variables and paths**

The direction of the slash for directory paths might vary in the documentation. Regardless of what you see in the documentation, follow these guidelines:

- **Linux** | **UNIX** Use a forward slash (/).
- **Windows** Use a backslash (\).

The names of environment variables are not always the same in Windows and AIX®. For example, %TEMP % in Windows is equivalent to \$TMPDIR in AIX or Linux.

For environment variables, follow these guidelines:

- **Linux** | **UNIX** Use `$variable`.
- **Windows** Use `%variable%`.

**Windows** If you are using the bash shell on a Windows system, you can use the AIX conventions.

### **Installation directory variable and paths for IBM Cloud App Management server**

`install_dir` is the installation directory for the IBM Cloud App Management server. For Red Hat Enterprise Linux operating systems, `/opt/ibm` is the default location.

### **Installation directory variable and paths for agents**

`install_dir` is the installation directory for the agents. The default location depends on the operating system:

- **Windows** `C:\IBM\APM`
- **Linux** `/opt/ibm/apm/agent`
- **AIX** `/opt/ibm/apm/agent`



---

## Chapter 4. Product overview

IBM Cloud App Management is a modern management platform, providing application-aware infrastructure monitoring and analysis for improved time to value. As a cloud-native application management platform, IBM Cloud App Management provides the right set of easy to use tools that meet the needs of development, operations, and site reliability engineering (SRE) teams. Tools needed to quickly find the root cause of an issue across a broad range of technologies, hybrid-cloud, and complex microservices architectures in all sorts of industries. IBM Cloud App Management delivers app-centric monitoring of microservices-based applications in addition to monitoring for traditional resources across enterprises.

IBM Cloud App Management includes IBM Cloud Event Management and both are included with the IBM Cloud Pak for Multicloud Management, running on Red Hat OpenShift.

IBM Cloud App Management gives site reliability engineers (SREs) a consistent monitoring method across the enterprise to any public or private cloud. Deploy in minutes, simplify your application management with increased flexibility, and deliver on different aspects of the application modernization journey.

With IBM Cloud App Management, you can simplify monitoring and incident management, which helps decrease time to resolution no matter how complex the hybrid microservices-based environment.

### **Support for application modernization**

Technology is becoming more integral to the success and competitive advantage of the business. Disruption is all around and companies need to respond or risk being displaced. Modernizing applications is imperative and monolithic applications are being transformed to containerized and microservices based architectures.

Organizations are moving from a waterfall style of development to practicing continuous delivery with DevOps and site reliability engineering.

On-premises data centers are being replaced by private and public clouds.

Your organization might use IBM Tivoli Monitoring (V6) and IBM Cloud APM (V8) to monitor your environment. Existing agents can run in dual mode with some agents that are connected to the Tivoli Enterprise Monitoring Server or Cloud APM server and others to the Cloud App Management server.


Whether you're modernizing your existing applications or building cloud native applications to adapt to new market opportunities, IBM Cloud App Management provides a path forward, bridging your investments in traditional workloads and modern cloud-based environments.

### **Enhances application resiliency**

Scalability and resiliency are built into this cloud-native app management solution. Cloud App Management can operate at a high scale and handle the dynamic nature of microservices-based applications and technologies.

Kubernetes views show service metrics from automatically normalized *golden signals* – latency, errors, traffic, and saturation – for quickly assessing the health of a service. Use these as early warning signals to get ahead of service impacts.

The golden signals and other visualization features, like the one hop topology to show the immediate dependency of the current service and the adjustable timeline with event markers to guide you, shorten your time to resolution.



### **Delivers a single management solution**

IBM Cloud App Management provides a single solution for developers, DevOps teams, and IT operations to manage their middleware and modern microservices-based business applications. They use the same centralized tool that is portable across the infrastructures.

Using IBM-provided webhooks, you can set up event feeds from competitive offerings to send data to. For example, you can set up an integration to receive notifications about jobs from Jenkins projects and set up another integration to receive alert information from Microsoft Azure.

You can also integrate with other products to send notifications and metrics. For example, you can set up an outgoing integration to send incident information to a GitHub repository as an issue.

### **User monitoring**

IBM Cloud App Management gives you the ability to measure availability, response time, and user satisfaction from a single control point across geographical locations.

### **Kubernetes monitoring**

Visually identify problems in Kubernetes clusters with dynamic view of nodes, pods, and containers correlated to activity in the cluster. Quickly see all resources that are contained in a cluster and their status. See the top nodes with highest usage to quickly take corrective action. Drill down into the individual node to view all the pods and containers within the node. With IBM Cloud App Management, you see the most important information that is affecting the Kubernetes services.

### **Dashboard, reports, and analytics**

Support DevOps with metrics history, trends, and forecasts. Query information about all the clusters that are connected.

### **Resource monitoring**

You can monitor application resources on multiple clusters. Collect metrics from infrastructure and get proactive alerts on threshold breaches.

### **Integrated event management**

IBM Cloud Event Management installed along with IBM Cloud App Management gives you the capability to automate management of incidents and events that are associated with resources and applications. You can visualize and manage multiple clusters, and consolidate the information from your monitoring systems to address problems. Events can indicate something that happened on an application, service, or another monitored object. All events that are related to a single application, or to a particular cluster, are correlated with an incident. Event Management can receive events from various monitoring sources, either on-premises or in the cloud.

## Offerings and features

IBM Cloud App Management contains the IBM Cloud App Management, Base and IBM Cloud App Management, Advanced offerings. These offerings manage traditional and Kubernetes resources and use advanced correlation and automation.

Click the following links to quickly access to the information that you want:

- [Offering components](#)
- [Product features and capabilities](#)

The offerings include the following components.

### Cloud App Management server

For more information, see [Chapter 9, “Installing IBM Cloud App Management - the options,”](#) on page 93.

### ICAM Agents

For more information, see [“Descriptions”](#) on page 54.

### ICAM Data Collectors

Contains: Kubernetes data collector, HMC agent, and runtime data collectors including Go data collector, J2SE data collector, Liberty data collector, Node.js data collector, Python data collector, and Ruby data collector. For more information, see [“Descriptions”](#) on page 54.

### Unified Agent

The Unified Agent is Unified Agent is a cost effective solution for development and maintenance. It integrates the open source technologies and has the capacity of collecting metrics, tracing, event, and so on. The Unified Agent provides a lightweight plug-in architecture, supports cloud native environment, and is easy to expand.

Use the Unified Agent to collect, process, aggregate, and write metrics to your Cloud App Management environment. It is based on Telegraf. For more information, see [Chapter 16, “Deploying Unified Agent,”](#) on page 645.

### IBM Cloud App Management Extension Pack

The IBM Cloud App Management Extension Pack extends Cloud App Management system monitoring to other environments. The extension pack starts with the SAP HANA Database agent, which provides usage data such as memory and CPU usage, database locks, and critical alerts. Database administrators can use this information that is collected by the SAP HANA Database agent to complete monitoring and other tasks such as responding to alerts. IBM Cloud App Management Extension Pack is available with both IBM Cloud App Management, Base and IBM Cloud App Management, Advanced offerings.

### IBM® Multicloud Manager Event components

IBM Cloud App Management offers a comprehensive cloud monitoring solution. You can visualize and monitor multiple clusters when you install IBM Cloud App Management in an IBM Multicloud Manager environment. The following components support IBM Cloud App Management when installed in an IBM Multicloud Manager environment.

- IBM Cloud App Management for Eventing Klusterlet Config
- IBM Cloud App Management for Eventing Klusterlet Config on AMD64
- IBM Cloud App Management for Eventing Server Side
- IBM Cloud App Management for Eventing Server Side on AMD64PLinux

### IBM Event Correlation for IBM Cloud App Management add-on

IBM Event Correlation for IBM Cloud App Management enables DevOps and IT operations teams to help resolve application, service, and infrastructure issues quickly by processing events from multiple third-party sources and also existing IBM infrastructure.

**Note:** To use the event source integration features (marked with an asterisk (\*) in the table) with IBM Cloud App Management, Advanced, you must order the IBM Event Correlation for IBM Cloud App Management add-on. Contact your IBM sales team for details about how to order this add-on.

For more information about part numbers and file names for the IBM Cloud App Management components, see “Part numbers” on page 71.

The following table lists the features and capabilities that IBM Cloud App Management can provide.

*Table 1. Features in each offering*

Feature	IBM Cloud App Management, Base	IBM Cloud App Management, Advanced	Links
Configure users and groups	✓	✓	<a href="#">How to</a>
View metering metrics	✓	✓	<a href="#">How to</a>
Create thresholds	✓	✓	<a href="#">How to</a>
Native event source integration	✓	✓	<a href="#">How to</a>
Event source integration with Datadog	—	✓*	<a href="#">How to</a>
Event source integration with New Relic Legacy	—	✓*	<a href="#">How to</a>
Event source integration with Amazon Web Services	—	✓*	<a href="#">How to</a>
Event source integration with Microsoft Azure	—	✓*	<a href="#">How to</a>
Event source integration with Netcool/OMNIBus	✓	✓	<a href="#">How to</a>
Event source integration with Jenkins	—	✓*	<a href="#">How to</a>
Event source integration with Pingdom	—	✓*	<a href="#">How to</a>
Event source integration with AppDynamics	—	✓*	<a href="#">How to</a>
Event source integration with Nagios XI	—	✓*	<a href="#">How to</a>
Event source integration with SolarWinds	—	✓*	<a href="#">How to</a>
Event source integration with Splunk Enterprise	—	✓*	<a href="#">How to</a>
Event source integration with Webhook	—	✓*	<a href="#">How to</a>
Event source integration with Logstash	—	✓*	<a href="#">How to</a>
Event source integration with Elasticsearch	—	✓*	<a href="#">How to</a>
Event source integration with Dynatrace Splunk	—	✓*	<a href="#">How to</a>
Event source integration with IBM Urban Code Deploy	—	✓*	<a href="#">How to</a>
Create event policies and runbooks for IBM Cloud App Management generated events	✓	✓	<a href="#">How to</a>
Create event policies and runbooks for external event sources	—	✓	<a href="#">How to</a>
Send incident details to Alert Notification	—	✓	<a href="#">How to</a>
Send incident details to Netcool/OMNIBus	✓	✓	<a href="#">How to</a>
Send incident details to Slack	✓	✓	<a href="#">How to</a>

Table 1. Features in each offering (continued)

Feature	IBM Cloud App Management, Base	IBM Cloud App Management, Advanced	Links
Send incident details to Webhook	—	✓	<a href="#">How to</a>
Send incident details to Microsoft teams	—	✓	<a href="#">How to</a>
Send incident details to Stride	—	✓	<a href="#">How to</a>
Send incident details to Service Now	—	✓	<a href="#">How to</a>
Send incident details to GitHub	—	✓	<a href="#">How to</a>
Send incident details to Watson™ Workspace	—	✓	<a href="#">How to</a>
View, investigate, and resolve incidents	✓	✓	<a href="#">How to</a>
Use the Resources view to visualize the metrics that are related to ICAM Agents and ICAM Data Collectors.	✓	✓	<a href="#">How to</a>
Use the Resources dashboard to visualize metrics gathered by the Unified Agent: <ul style="list-style-type: none"> <li>• UA plug-in for Jaeger and Zipkin</li> <li>• UA plug-in for NGINX</li> <li>• UA plug-in for Redis</li> <li>• UA plug-in for IBM API Connect(APIC)</li> <li>• UA plug-in for IBM App Connect Enterprise(ACE)</li> <li>• UA plug-in for IBM MQ</li> <li>• UA plug-in for DEM</li> <li>• UA plug-in for OpenShift</li> </ul>	✓	✓	<a href="#">How to</a>
Go back in time to visualize the state of each Kubernetes resource layer at the time an event was fired	—	✓	<a href="#">How to</a>
Create synthetic tests and monitor response time and availability for your Rest API websites.	—	✓	<a href="#">How to</a>

Table 1. Features in each offering (continued)

Feature	IBM Cloud App Management, Base	IBM Cloud App Management, Advanced	Links
<p><b>Digital Experience Monitoring (DEM)</b>            Digital experience monitoring (DEM) can be enabled in IBM Cloud App Management to monitor web-based resources and real user experience. It can discover and track traffic, user behavior, and other metrics to help analyze the application performance and usability.</p> <p>IBM Cloud App Management provides two ways to enable DEM. You can enable DEM for Liberty data collector, and install the DEM plug-in for HTTP Server to monitor IBM HTTP Server and Apache HTTP Server.</p>	✓	✓	<a href="#">How to</a>
<p><b>Transaction Tracking</b>            The transaction tracking feature enables topology views and instance level transaction monitoring. By distributed tracking infrastructure, transaction tracking can detect bottleneck issues including latency problems and errors, and filter or sort traces based on application. Transaction tracking can also filter views based on length of trace, timestamp, interactions, errors and transaction comparisons. For more information, see <a href="#">“Transaction tracking”</a> on page 780.</p>	✓	✓	<a href="#">How to</a>

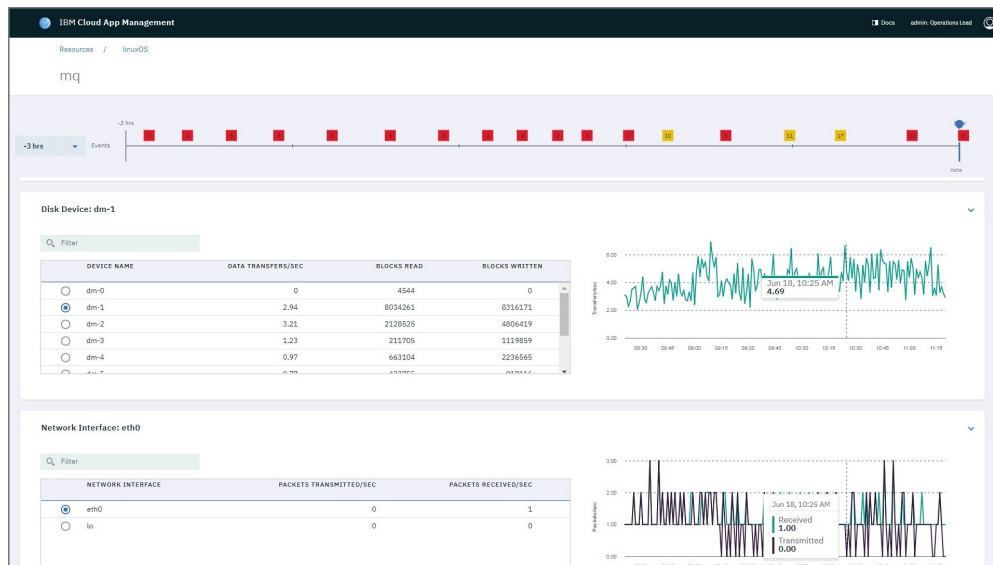
## User interface

The Cloud App Management console is the user interface for monitoring your mission-critical applications.

Based on the award winning design for IBM Cloud Event Management, IBM Cloud App Management is interactive and scalable:

- In the **Resources** dashboards, you can select other metrics that you'd like to see and compare, use the time slider to adjust the dates shown, and use the time selector to adjust the time range from the past 3 hours to the past month.
- You can quickly sort through and filter lists and tables to show only what you're interested in.

Use the console to check the status of your applications and respond to incidents. The dashboards simplify problem identification with the incident management capability and dashboard navigation that takes you from a view of application status to code level detail. You have visibility into source code problems at the exact moment of an issue.



Take a look at the usage scenarios to learn more about what you can do in the Cloud App Management console.

## Getting started: Manage dynamic application and infrastructure environments

As a developer or IT operator, you want to be able to quickly isolate and focus on issues affecting your application or the environment that is hosting your application. Follow this scenario to generate some sample incidents and learn about the incident queue in the Cloud App Management console.

Using the Cloud App Management console to manage dynamic application and infrastructure environments is beneficial:

- Ensure the performance of applications and application infrastructure, with quick time to value, while driving down your IT management total cost of ownership.
- IBM Cloud App Management is a cloud native management platform which provides unique insights to manage your complex application environment – on premises, private or public cloud environment, or in a hybrid environment covering any combination

After your Cloud App Management environment is set up, thresholds are activated to test for resource issues such as a database failure and slow response time. When the conditions of a threshold are true, an event is opened and an incident is generated.

Cloud App Management takes open events and correlates and groups them into incidents to reduce the noise. Take Cloud App Management for a test drive by generating some sample incidents, then viewing the incident queue.

### Generate sample incidents


1. Select **Administration** from the menu bar.
2. Click the **Integrations** hexagon or its Information icon.



3. Click **Generate sample incidents** in the **Information** box or click **Generate** in the **Sample Events** box.
4. When you see the **Success notification** message, click **Go to incidents**.

## View the incident queue


The incident queue presents the sample incidents that you just generated sorted by highest priority and then by most recently changed, so you can easily see the most urgent incidents.



The incident queue has some different characteristics depending on your how your Cloud App Management environment was configured and whether setup is complete.

### Are your agents deployed?

Along with the sample incidents, you see incidents generated by any events from your deployed Cloud App Management agents or integrated Cloud APM V8.1.4 agents or Tivoli® Monitoring agents.



Type samp in the Search box to see only the sample incidents; or try exp to see only those with "experiencing" in the description.

### Do you have incident policies?

Incident policies perform actions against an incident. Cloud App Management has built-in incident policies that assign a priority to the incident based on the event severity.



Order	Name	Actions	Last modified	Last run	Enabled
1	Set Priority 5 Author SYSTEM	Priority	Jun 4, 2018   11:07:52 AM PDT	Jun 7, 2018   1:42:05 PM PDT	On
2	Set Priority 4 Author SYSTEM	Priority	Jun 4, 2018   11:07:52 AM PDT	Jun 7, 2018   1:42:05 PM PDT	On
3	Set Priority 3 Author SYSTEM	Priority	Jun 4, 2018   11:07:52 AM PDT	Jun 7, 2018   1:42:05 PM PDT	On
4	Set Priority 2 Author SYSTEM	Priority	Jun 4, 2018   11:07:52 AM PDT	Jun 7, 2018   1:42:05 PM PDT	On
5	Set Priority 1 Author SYSTEM	Priority	Jun 4, 2018   11:07:52 AM PDT	Jun 7, 2018   1:42:04 PM PDT	On
6	policy1 Author admin	Assign and notify	Jun 5, 2018   6:29:53 PM PDT	Jun 7, 2018   1:57:02 PM PDT	On

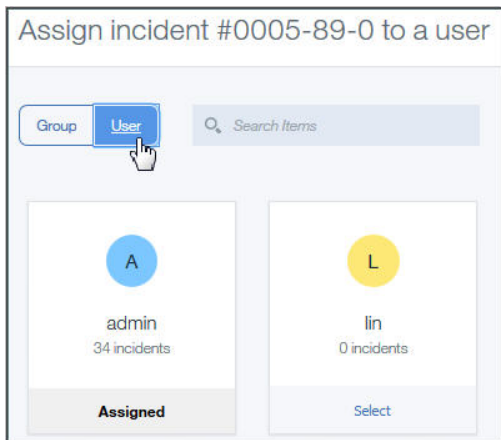
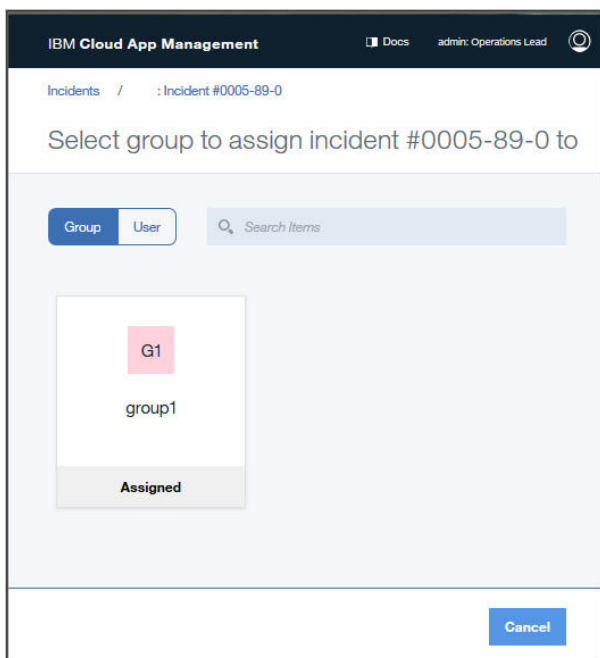
Any policies that were defined for your environment can affect incidents. For example, if you have a policy that assigns all priority 4 incidents to a particular user, you'll see the two sample incidents assigned to that user:

### Have users or groups been added?

If Cloud App Management users or groups have been added, they can be assigned to incidents and can collaborate on resolving incidents. Incident policies can be defined to automatically assign one or more users or a group to all incidents or incidents with certain characteristics such as a specified type of incident.

Do you see an **Unassigned** sample incident? You can assign it in one of these ways:

- Drag and drop a user or group from the sidebar over an unassigned sample incident.
- Click **Incident actions** on an unassigned sample incident and select **Assign**. In the assignment page that opens, click **Select** in a user or group box.



The status changes to **Assigned**.

### Let the sample incidents expire

The sample incidents give you an opportunity to test drive Cloud App Management incidents even if you haven't finished setting up your environment for resource monitoring. They expire in an hour, and you can generate sample incidents again at any time. Any incidents that were resolved are removed from the queue two minutes later.

## Getting started: Collaborate to rapidly resolve problems

As the team lead, you must quickly assess and prioritize incidents as they arrive. You assign incidents to other users or work on them yourself to keep the work evenly distributed. Follow this scenario to learn more about incident research and discovery in the Cloud App Management console.

Using the Cloud App Management console to collaborate to rapidly resolve problems is beneficial:

- Collaborate with team members to quickly and effectively handle incidents and problem diagnosis within your application environment, reducing impacts to users and to your business.
- Events are de-duplicated and visualized to give you a unified view of the incidents impacting your application environment, allowing you to organize, prioritize, assign, notify, and diagnose in context – leading to rapid resolution.

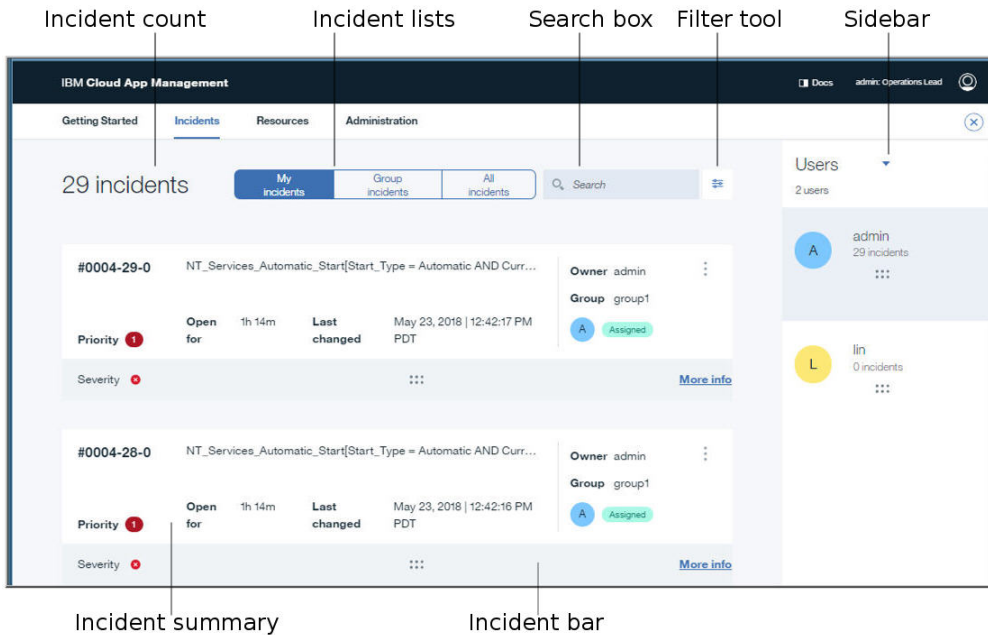
In this scenario, you'll assign one incident and work on another to find the root cause using the resource dashboards.

## Open your incident queue

Whether you are notified of an incident by email or you are already in the Cloud App Management console, you can get to your incident queue in one of these ways:

- Click the link in the email to go directly to incident entry in your queue.
- In the landing page after logging in to the console, select **Go to my incidents**.
- Select **My incidents** from the **Getting Started** page.
- Select **Incidents** from the menu bar.


A summary view of each incident that is assigned to you is displayed.



## Find the incidents that aren't being worked on

Initially, you see the incidents that are assigned to you, but you want to also see them for the entire team so you click **Group incidents**.

You can use the search box to find incidents by their ID number or summary description. In this case, you want to see all the higher priority incidents that aren't being worked on: Click **Filter** and select the **Assigned** status check box, then the **Priority 1** and **Priority 2** check boxes.




Two Priority 1 incidents are shown and both are assigned to another user. You read from the incident description that one incident is for high CPU usage and the other is for an inactive database.

### Reassign an incident to another user

The sidebar shows all the users in the group. You see that Steven has no incidents, so you reassign the Priority 1 database incident by dragging the **⋮** grippy from the Incident bar and dropping it on his name in the sidebar. Now you have only the CPU incident to deal with.

### Take ownership of an incident

You decide to take the CPU incident yourself because the assignee currently has 8 incidents in their queue. Click **⋮ Incident actions** and select **In progress** to show that you are working on it.

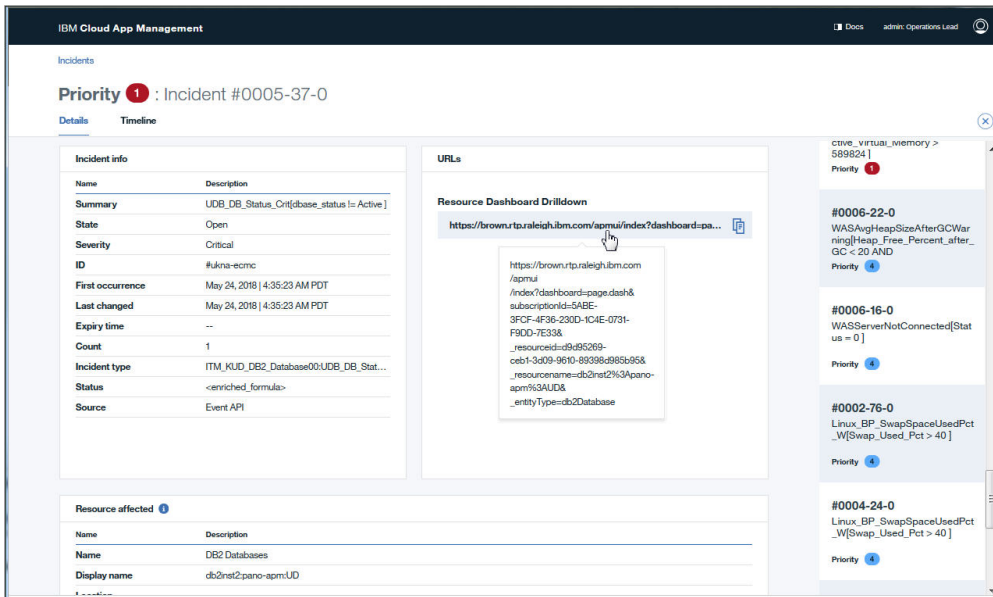


### Get more information about the incident

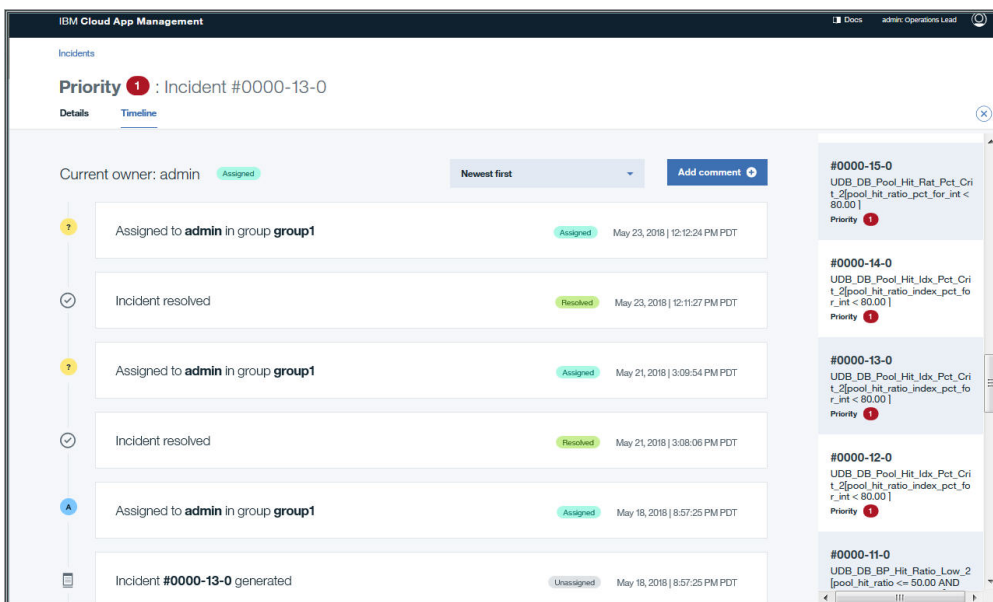
You click **More info** and a tabbed page opens:

- The **Details** tab has information about the event and incident. The **First occurrence** and **Last changed** fields in the **Incident info** area tell when the threshold (or situation from your integrated Tivoli Monitoring agent) was first breached and the most recent change. Focus on this time range when viewing the dashboard metrics to help locate the cause of the event. The **Count** field shows

the number of *deduplicated events*, which are multiple occurrences of the same event.



- The **Timeline** tab shows the changes in incident status since it was first opened. You can add a comment about the incident here.




### Open the resource instance dashboard

You have a link in the **Details** tab that opens the resource instance dashboard related to the incident.

The time slider shows pins, which are dropped where events occurred, so you can easily research what caused them. Click a pin to see the metrics at the point in time when the event surfaced. You can also drag the slider to see before and after the event.

Scroll down to see all the metrics.



## Resolve the incident

Using the dashboard tools, you're able to find a pattern of high CPU usage when certain applications are running and can take action to resolve the incident.

You navigate to the incident resolution page where you can add notes about the cause of the high CPU, what you did to solve the problem, and suggest actions to prevent a recurrence such as running the applications at off-peak hours or reconfiguring them to use less memory:

1. Click the **Resources** breadcrumb to return to the front page, then click **Incidents** from the menu bar.
2. To resolve the incident, you could select **Incident actions > Resolve** from your incident queue. But because you also want to enter a comment, click **More info**, select the incident from the sidebar to open the incident resolution page, and select the **Resolve** button.
3. Click **Add comment**, then write down the actions you took to solve the problem and your suggestions for avoiding a repeat of the same issue.

## Getting Started: Proactively manage the health of your application environment – regardless of size

You're the operations lead and want to automate some incident handling by adding a new policy. Follow this scenario to learn more about incident policies and user profiles and how they are manifested in the incident queue.


Using Cloud App Management console to Proactively manage the health of your application environment – regardless of size is beneficial:

- Proactively manage your application environment by finding and fixing application problems BEFORE your users are impacted.
- IBM Cloud App Management is a highly resilient solution, built to handle the dynamic scale of your application environment, with a design that lets you quickly search and filter to focus on the resources and their relationships that matter most to you. Designed with automation in mind, API's are used to provide hands off administration, making proactive management easier and helping lower cost of ownership.

## Review the incident policies

In the first Getting Started scenario, you generated sample incidents and saw how they were prioritized based on the built-in policies and any custom policies in your environment. Familiarize yourself with the policy options.

1. From the Cloud APM console menu bar, select **Administration** > **Policies** to open the incident policies page.



2. Click **Actions menu** to see the options for moving the policy up or down the list. If a policy has a conflicting rule with one that comes earlier in the list, the rule of the policy that comes after overrides the earlier one.

<input type="checkbox"/>	Order	Name	Actions	Last modified	Last run	Enabled	
<input type="checkbox"/>	1	Set Priority 5 Author SYSTEM	Priority	Jun 4, 2018   11:07:52 AM PDT	Jun 10, 2018   9:41:48 PM PDT	<input checked="" type="checkbox"/> On	<b>Edit</b>
<input type="checkbox"/>	2	Set Priority 4 Author SYSTEM	Priority	Jun 4, 2018   11:07:52 AM PDT	Jun 10, 2018   9:41:48 PM PDT	<input type="checkbox"/>	Move up
<input type="checkbox"/>	3	Set Priority 3 Author SYSTEM	Priority	Jun 4, 2018   11:07:52 AM PDT	Jun 10, 2018   9:41:48 PM PDT	<input type="checkbox"/>	Move down
<input type="checkbox"/>	4	Set Priority 2 Author SYSTEM	Priority	Jun 4, 2018   11:07:52 AM PDT	Jun 10, 2018   9:41:48 PM PDT	<input type="checkbox"/>	Move to top
							Move to bottom
							Delete

3. Select **Actions menu** > **Edit** for the Set Priority 5 policy:

**IBM Cloud App Management** Docs admin: Operations Lead

Administration / Policies

## Edit incident policy

What's an incident policy? It selects incidents by using conditions — either on the attributes of incidents, or on the events correlated to incidents. It also applies actions to the matching incidents. Actions include sending notifications about incidents, assigning incidents to users and groups, and assigning runbooks to incidents.

Author: SYSTEM  
 Created: Jun 4, 2018 | 11:07:52 AM PDT  
 Last updated: Jun 4, 2018 | 11:07:52 AM PDT  
 Last updated by: SYSTEM  
 Last applied: Jun 10, 2018 | 9:41:48 PM PDT

### 1. Details

**Policy name:** Set Priority 5

**Description:** Default policy to identify priority 5 incidents. Change the policy conditions to match your requirements.

### 2. Incidents

All incidents  Specify conditions

> Information and examples

**Conditions:**

Incident has the following attributes

Attribute: Priority is higher than or equal to Priority 5

AND

Describe the events that the incident contains:

Condition 1: Incident contains the following events

Attribute: Severity is less than Information

AND

Find out what incidents would have matched this condition in the past 3 Days **Test**

### 3. Action

Assign and notify

Set priority

Set the priority for the incidents described above.

Priority 5

### 4. Enable

Off  On

Cancel Save

This policy assigns the incident priority for information events.

- a. The **2. Incidents** option is set to **Specify conditions**.
- b. The condition attribute is **Priority** and applies to Priority 5 (or higher) incidents. This incident policy has only one condition and it is based on priority, but you can apply more conditions to a policy based on priority, assigning user, or when the incident was last changed.




- c. The incident is for events of severity information or lower. You can have multiple event attributes such as for a specific host name. Some commonly used event types are listed in **Add predefined conditions**.
  - d. **3. Action** sets the incident to Priority 5. The **Assign and notify** check box wasn't selected for this policy, but you could also have the incident assigned automatically to any combination of users, groups, and integrations.
4. Click **Cancel** to close the Set Priority 5 definition.

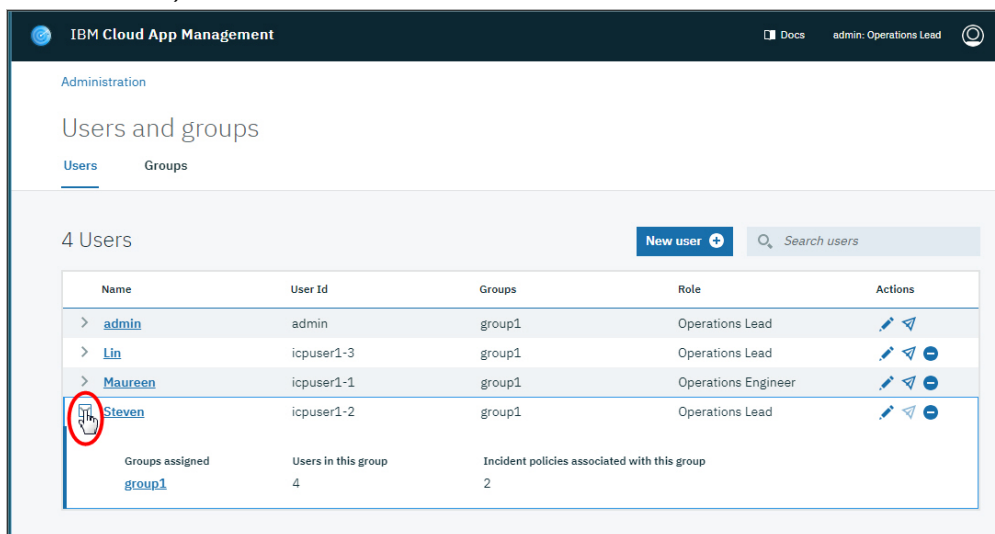
### Review user profiles

You saw how incident policies can be set to automatically assign an incident to one or more users, groups, or both. Take a look at the profile options for users.

1. From the Cloud APM console menu bar, select **Administration > Users and Groups** to open the **Users** tab.



The Users tab lists all the users by their name, IBM Cloud Private user ID, the group they belong to, their role and notification status (a dimmed  **Notify** icon means the user hasn't verified their email address).



Your environment has an IBM Cloud Private admin ID, which cannot be deleted.

2. If a user ID has a group assigned, you can click the twisty to expand the entry to see how many users are in the group and how many incident policies are assigned to this group.


### Add a group

Because you're the operations lead, you can create a new group and assign yourself as the owner:

1. After reviewing user IDs, select the **Group** tab. Just as with users, you can click a twisty to expand an entry and see the group members and which policies they are associated with.
2. Click **New Group**.
3. You name the group database team and add the database administrators to the membership.
4. By assigning yourself as the owner, you can manage the group membership.
5. After you save the new group, return to the **Administration** front page.

### Create an incident policy

1. Select **Administration > Policies**.



2. Click **Create incident policy**.
3. Name the policy Database issues.
4. Specify the condition **Priority is higher than Priority 4** with **Resource type is DB2 Instances**.
5. Select the **Assign and notify** check box and assign to the database team and to the database team lead Steven.

## Getting started: Accelerate your transition to the cloud with DevOps

What do you do when you find out about a problem not from an incident but from a help ticket? Follow this scenario and learn some proactive measures you can take to avoid future problems.


Using the Cloud APM console to Accelerate your transition to the cloud with DevOps is beneficial:

- Developers and IT Operations are working together to deliver innovation at speed and scale, leveraging cloud native technologies such as microservices, containers, Kubernetes and DevOps. Successful enterprises are adopting these new technologies for new cloud-native applications and for modernizing their existing ones to deliver business agility.
- IBM Cloud App Management provides a single solution for Developers and Operations teams, which seamlessly integrates into your DevOps practices and toolchain.

In the previous scenarios, you learned about the incident queue and the features for managing and handling incidents. In this scenario, the ITops team notice some peculiar behavior that they'd like to monitor. You'll define a threshold and use the resource dashboards to help you fine tune the formula.

### Open the Threshold Management page

In the Cloud APM console, click **Administration > Threshold**.



In the Thresholds front page, you see a list of thresholds ordered by highest severity. For each threshold, you can see the resource it monitors, whether it is read-only (such as predefined thresholds) or editable, and whether the state is enabled and or disabled.

The screenshot shows the IBM Cloud App Management console interface. At the top, there's a dark blue header with the IBM logo and 'IBM Cloud App Management' text, and a 'Docs' link on the right. Below the header, the navigation menu includes 'Administration'. The main content area is titled 'Threshold management'. In the top right of this area, there is a 'Create' button with a plus icon, circled in red, and a search bar labeled 'Filter threshold'. Below this is a table with the following columns: Name, Severity, Assigned to, Permissions, and State. The table contains two rows of data:

<input type="checkbox"/>	Name	Severity	Assigned to	Permissions	State
<input type="checkbox"/>	PoorDevilNtwkPktsRcvd	Critical	Linux Systems	Editable	Enabled
<input type="checkbox"/>	HuggyBearNtwkPacketsRcv	Warning	Linux Systems	Editable	Enabled

### Create a threshold

You create a threshold. ITops told you they were having disk file space issues on their Linux systems, so you create a threshold to monitor the percentage of time spent in read operations:

1. Click **Create** to define a new threshold.
2. For the condition, you want to open a warning event for a Linux Systems resource if the Disk IO Disk Read Percent is greater than or equal ( $\geq$ ) to 80%.
3. You name the threshold policy Check\_disk\_reads with a helpful description like, Warning event for disk read time of 80% or more.
4. You want to monitor specific resources rather than all Linux-type resources: For the Assign to resources step, you select **Individual instances** and, in the **Filter instances** list, you select the two resources that ITOps requested.

5. After you click **Save and finish**, the **Threshold Management** page shows the new threshold in the list and your threshold is tested every few seconds.

<input type="checkbox"/>	Name	Severity	Assigned to	Permissions <span>⬇</span>	State
<input type="checkbox"/>	Check_disk_reads	<span style="color: yellow;">●</span> Warning	Linux Systems	Editable	Enabled


### Review the Resources dashboard

Find out if the threshold you created is opening events where they're needed. Some thresholds might need fine tuning.

Click **Administration** to return to the front page, then click **Resources** from the menu bar. The most numerous resource types in your environment are displayed.

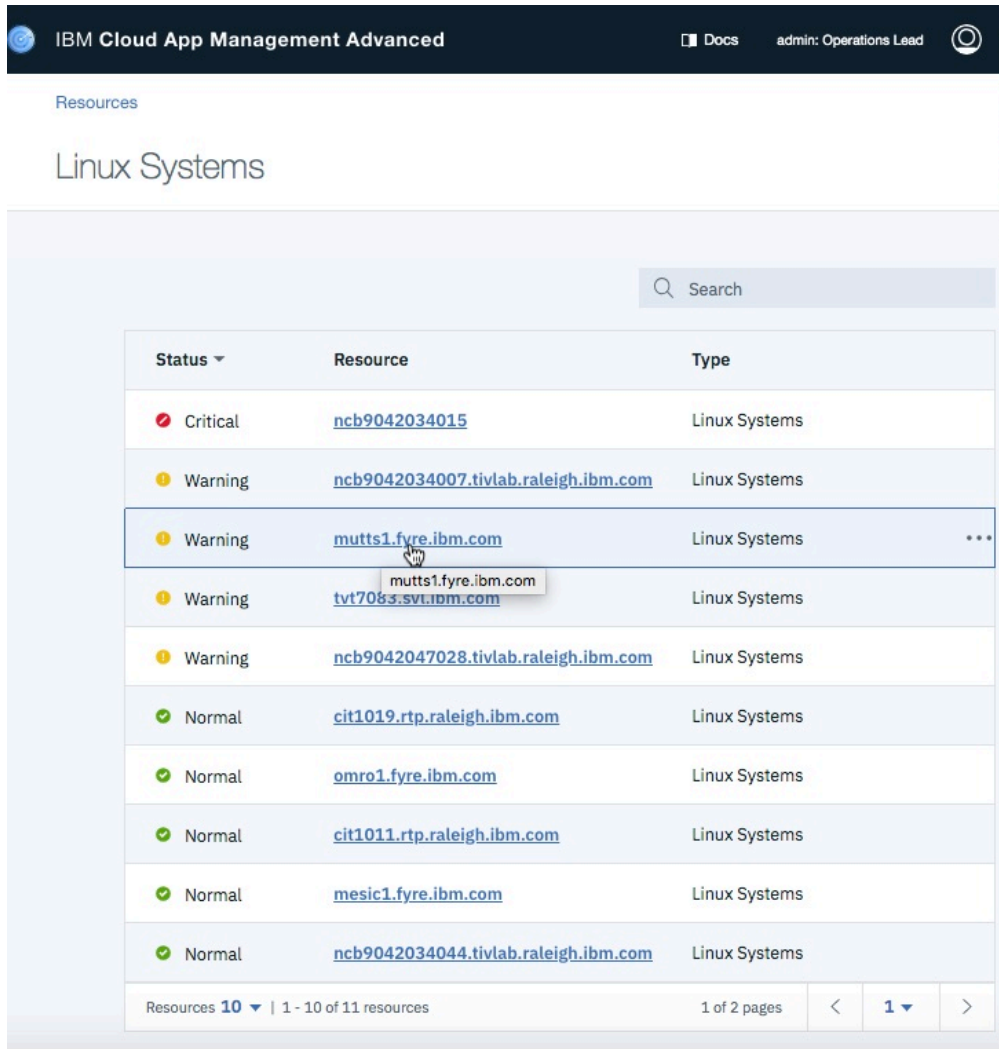
### All resources

Type `linux` in the search box to quickly find the Linux Systems resource type, and click **View Resources**.



### Resource type


After you select Linux Systems, the resource instances are displayed in a list sorted by its status.



## Resource instance

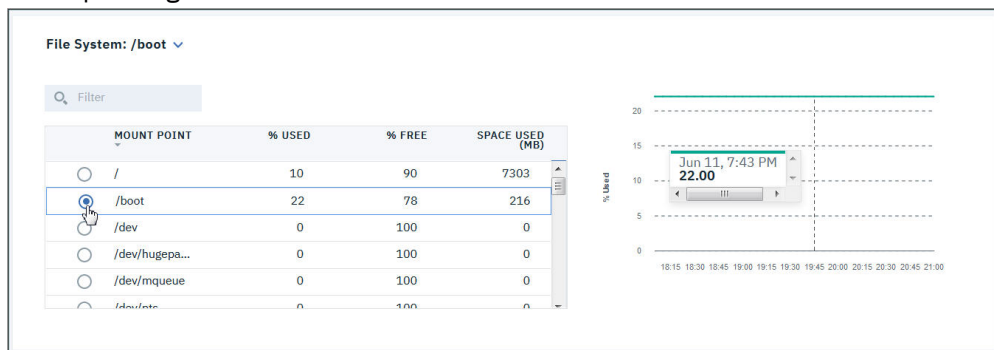
Select a link to open the dashboard for that instance. The instance dashboard is displayed with metrics from the past 3 hours by default. You can adjust the time span to show up to the past 32 days of saved samples from the data provider.

The time slider has dropped pins for events that occurred within the time range. Drag the timeline slider or click another point along the slider to dynamically update the metrics with the values from the time slot. You check each of the dropped pins, including the values before and after each event to see if you can find a pattern.




The Linux OS instance dashboard (also UNIX OS and Windows OS dashboards) plots metrics for the system characteristics. Dashboard sections with multiple metric views are synchronized:

- Click along the x-axis of the **Aggregate CPU Usage** or **Memory Usage** line chart to read the time stamp and value on both charts.
- Select a mount point in the **File System** table to see the percentage used over time in the corresponding line chart.



- Select a device name in the **Disk Device** table to see the transfers per second for that device in the corresponding line chart. You can sort the table by clicking a column, and filter it by entering a value in the filter box.




- Select the network interface type in the **Network Interface** table to see the number of packets transmitted and received per second plotted on the line chart.



The **System Information**, **Resource Properties**, and **Related Resources** tables have no counterpart. For these and the other dashboard sections, you can use the **▼ Collapse** and **► Expand** twisty to show only what's of interest to you.

### What if you want to see other metrics?


There's another metric you'd like to check: page outs, which might indicate memory issues. Scroll down to the **Custom Metrics** section, expand the view, and select System Pages Paged Out Per Second from the drop-down list. The list shows the metrics that are available for selection from the Linux data provider. You can add other metrics, but this metric and your analysis of the metrics around the event times is enough to tell you that the threshold you created needs a minor adjustment.



## Launch Threshold Management from Resources

Rather than navigating back to the threshold editor, you can launch in context from the resource dashboard:

1. Click the **Linux Systems** breadcrumb to return to the resource page for Linux Systems.



2. Click **Options > Thresholds** to open the Threshold Management page with the thresholds list filtered to show only those that are assigned to the resource instance.
3. Find the threshold that you created earlier and click **Options > Edit**.
4. Change the Disk Read Percent value from 80 to 75.
5. You tested the threshold on one resource and now want to disseminate it to all your Linux systems, so you change your selection in the **Assign to resources** section to Resource group.
6. Select **Save and finish** to see the edited threshold assigned to Linux Systems.

## Getting started: Performing SRE functions

How can Cloud App Management help simplify monitoring and quicken time to resolution? Let's look at a typical scenario to see how:

### Check the health of a business critical application

The Stock Trader application is a major financial earner for our company. The application has been refactored into microservices that are running both on premises and in the cloud. In production, the system of record is a shared pool of Db2 instances. We want to make sure this latest software update doesn't cause any performance issues.

In the Cloud App Management console, we navigate to the Kubernetes Service resources for the Trader application:

1. Click the **Resources** tab.
2. Search and locate the Kubernetes Service resource group. Display the list of resources in this group.
3. Filter the list to find the Trader-service resources.
4. From the multiple Trader-services resources that display, select the front end HTTP service that provides the interface to the user's browser.

The Service dashboard provides insight into how the service is performing based on the golden signals.


### Observe the golden signals

The golden signals are shown in the Latency, Error, Traffic, and Saturation line charts. We notice that the error rate is starting to increase. We need to quickly resolve these errors before they negatively affect our users. Looking at the causes – traffic and saturation – we quickly determine that Trader is not the cause of the increased latency.

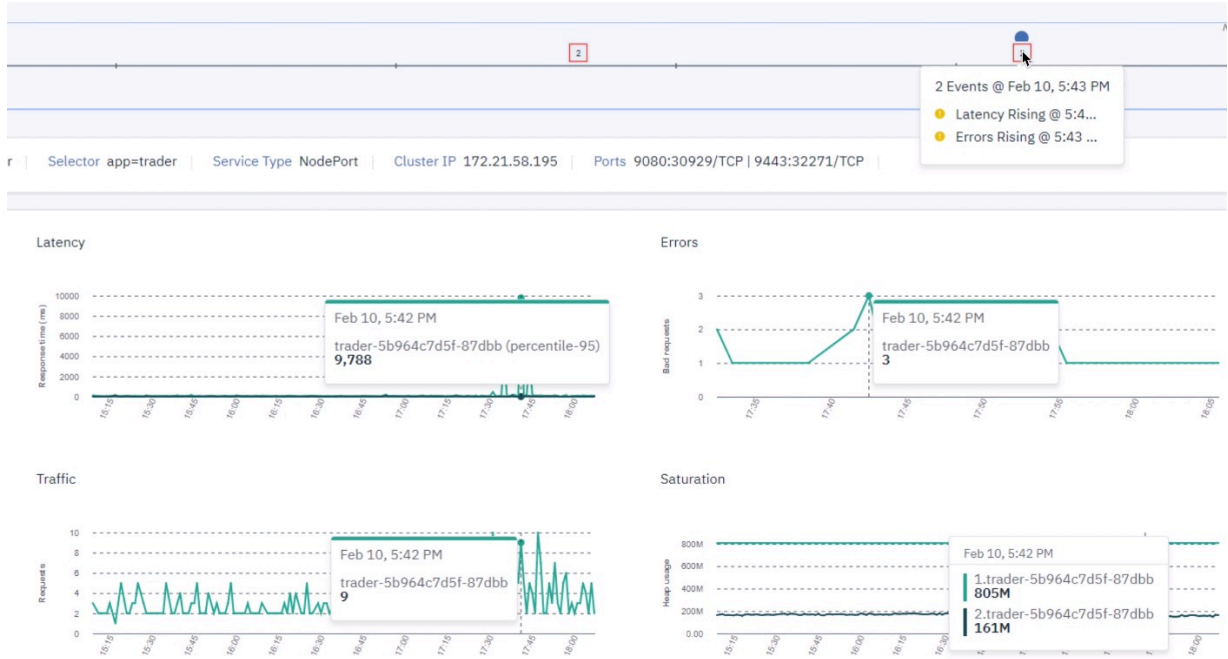
Typically, we would want to see if the problem is associated with a specific request type (or more) to determine where the root cause of the problem resides. We want to filter the signals by request to see which were impacted, and we notice that the latency increase is consistent across all request types; similarly for errors.

### Find the dependency with the 1-hop view

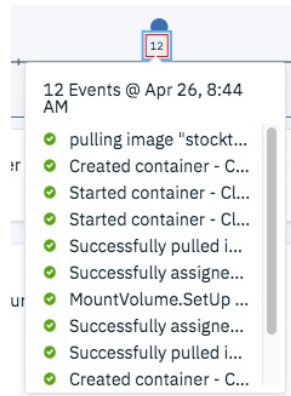
From the service 1-hop view, we see that Trader depends only on Portfolio. In this case, all requests go through Portfolio.



We click on the Portfolio service to open its dashboard. We don't see the same latency as Trader, however we see that the errors are still present. We see that Portfolio is Saturated, thus causing the errors.



We also confirm that the event from the recent code push is showing up in the timeline and look for a change in metrics at this point in time.




To gain some more insight, we use the service deployment view to navigate to one of the pods that implement this service. We can see that Portfolio is approaching the memory limit, at which time it will be killed and restarted by Kubernetes. Given that there is only 1 instance of Portfolio deployed, the 503 errors seen in Trader would be the result of requests issued while Portfolio was fully saturated.

Drilling down into the container view, we quickly see that Portfolio is exhausting Heap, causing the errors, the slow down, and what would be the eventual restart of the pod. Knowing that there was a recent deployment – and it is shown in the time slider as well – we determine that this update must have introduced a memory leak in Portfolio.



## JVM Garbage Collections

### Heap Used (KB)



To temporarily resolve this issue, we use a runbook to deploy another instance of Portfolio, adding capacity to reduce portfolio service saturation. Having another instance (or more) also alleviates the 503 errors because other instances can handle requests while another instance is being restarted.

IBM Cloud App Management Advanced

Getting started Incidents Synthetic tests Cloud resources Resources Administration

Integrations Configured

Runbooks Configured

Policies Configured

Thresholds Configured

Runbooks

- IBM Cloud Private - Scale a Deployment
- IBM Cloud Private - Set up the Calico CLI
- IBM Cloud Private - Set up the Helm CLI
- IBM Cloud Private - Setup LDAP
- IBM Cloud Private - Show Helm CLI
- IBM Cloud Private - View a Cluster
- IBM Cloud Private - View the Values
- Redeploy POD
- Reset an SA MP cluster application
- Restart DIZ instance
- Restart IBM Tivoli Monitoring Agent
- Restart Linux process
- Restart TADDM
- Scale Kubernetes Deployment

Execute runbook

2 parameters

DEPLOYMENT: portfolio

REPLICAS: 3

Procedure: 1 step, approximate time: a few seconds

STEP 1: Scale a Kubernetes deployment up or down. Information required: deployment name and number of running pods desired.

Scale Deployment

[kubectl scale deployment portfolio --replicas=3]

Successful

Complete

Provide feedback

Rate It: ☆☆☆☆

Comments

Given that this is a customer impacting problem, we hold a postmortem with the development team to understand how the memory leak was introduced and how it managed to find its way into production without being caught by automated testing in the pipeline.

## Incident Resolution Flow for a Kubernetes App

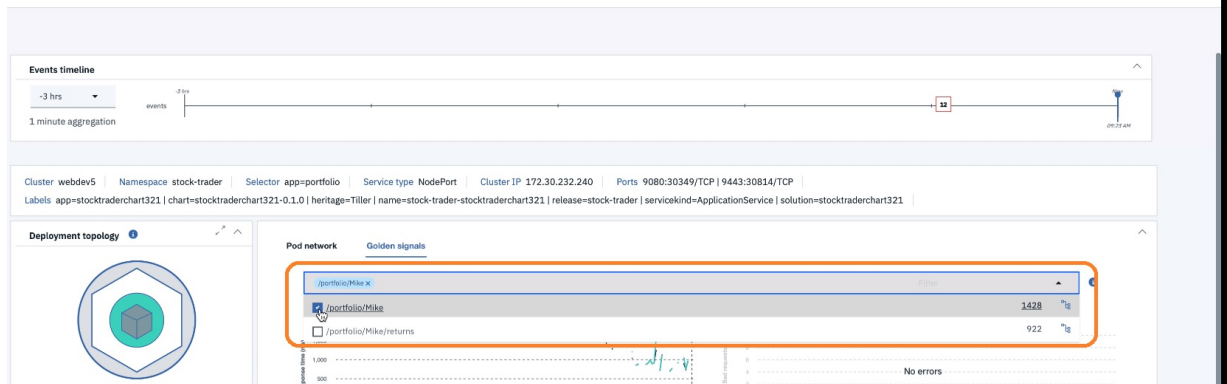
Resolving an incident in a Kubernetes service in IBM Cloud App Management. In this example, John (an SRE) is notified by an incident that is created with high latency with the stock trader service. John restores the service by creating a hypothesis and following it through to determine whether he isolated the problem. If the problem is not resolved, John creates another hypothesis. This example shows how this process is accomplished in IBM Cloud App Management.

### Is this latency problem affecting all of the service request types or just a subset?

To check this hypothesis, John filters the latency golden signal by request type. Examples of request types include: /view portfolio, /buy stock, or /sell stock. If just one request type is having latency

problems, John knows that the problem is localized to that one request type. He would start looking into the logic and logs to see what the service is experiencing. If it is all request types, then John is more likely to suspect a broader issue like infrastructure, network issues, or global dependencies for the service.

Resources / Kubernetes Service  
portfolio-service



### Are the service's dependencies affecting the latency?

To check this hypothesis, John looks at all the 1-hop away dependencies to see whether they are causing the slowness. He does this by looking at the Service dependencies widget and looking for any red services to the right, or downstream from his service. If so, John isolated the problem to another service and not his own. He would then move the focus to working on the dependency. If someone else owns the dependency, John can follow the established process to resolve with the owner. If the troubled dependency is another service John owns, he can select it to refocus IBM Cloud App Management on that service and start the hypothesis flow from the start.

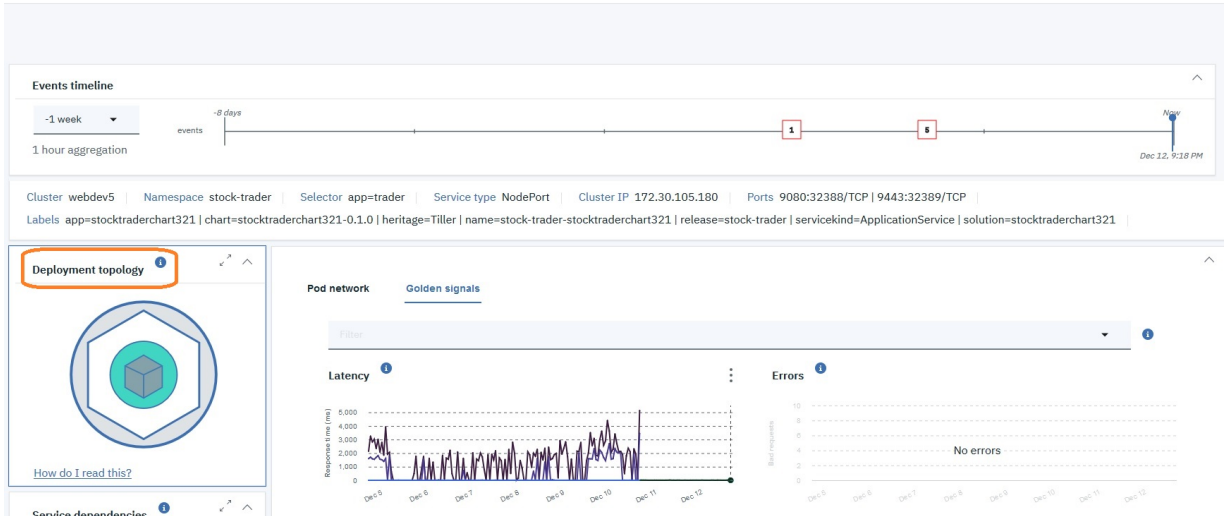
trader-service



### Is Kubernetes infrastructure impacting the service's latency?

To check this hypothesis, John looks at the Kubernetes deployment topology view to see whether there is any noisy neighbor problem by looking for red in the containers, pods, or nodes. If so, John knows that the infrastructure is the problem and not the service. To further isolate the problem, John can use the time slider to see what changed in the infrastructure to be causing the noisy neighbor problem. Also, he can look for Kubernetes infrastructure events on the timeline to highlight where the infrastructure problems occurred. One example of infrastructure impacting a service is over-consumption of compute resources. Kubernetes helps to protect from this issue if the best practices are followed.

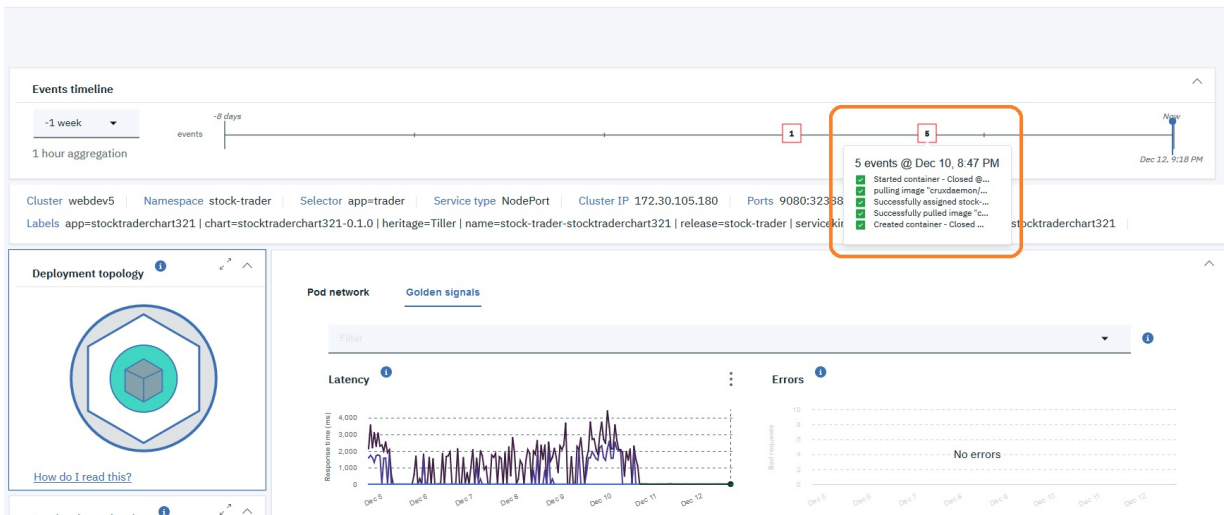
## trader-service



### Is this latency problem caused by new code checked in to the service?

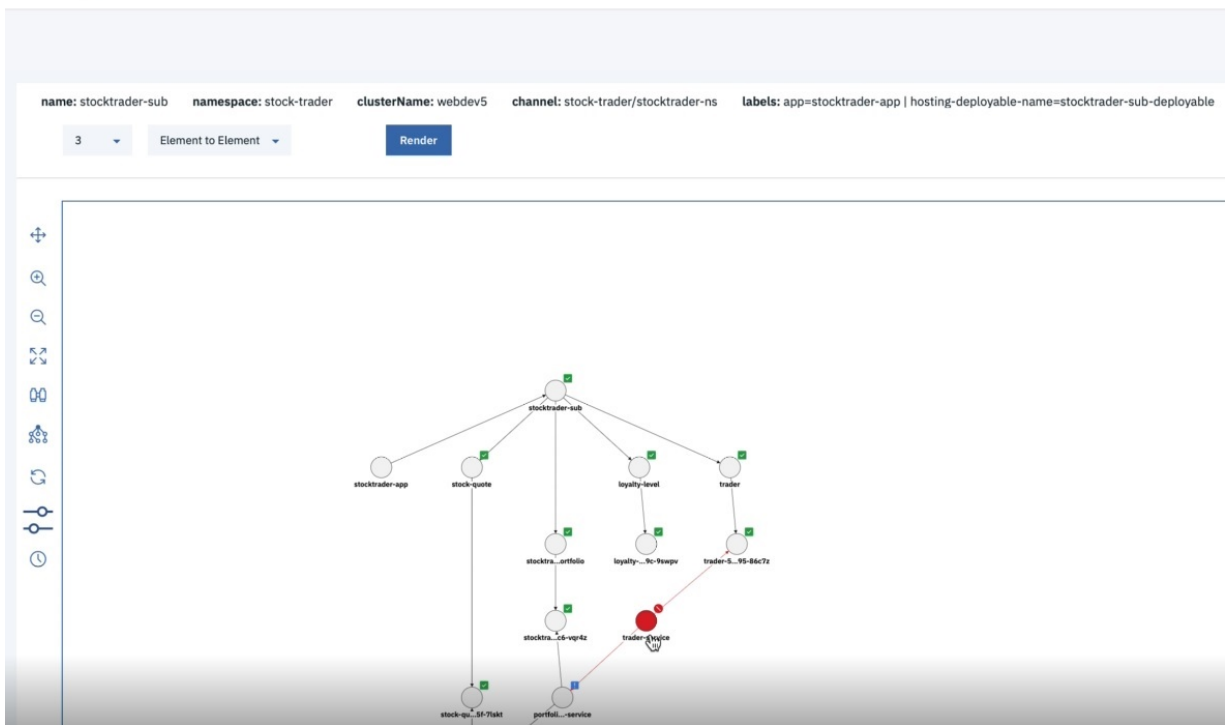
To check this hypothesis, John can go through the event timeline to see whether any CI/CD events are seen. John can check the latency before and after the change to see whether latency changes around the time of a new code deployment.

## trader-service



### Is this latency problem affecting the service caused by a network issue?

To check this hypothesis, John can select a wider topology view by expanding the Service dependencies view. This expanded view includes broader elements in the system including network elements on the topology. He looks for network events on the most problematic network area, for example, the gateway and load balancer, which connect users to the service. Examples of impacts to the service include: overloaded load balancers, slow LDAP servers, and DNS issues. If John has found an issue in the broader system, he can follow the established process to resolve them with the appropriate team.

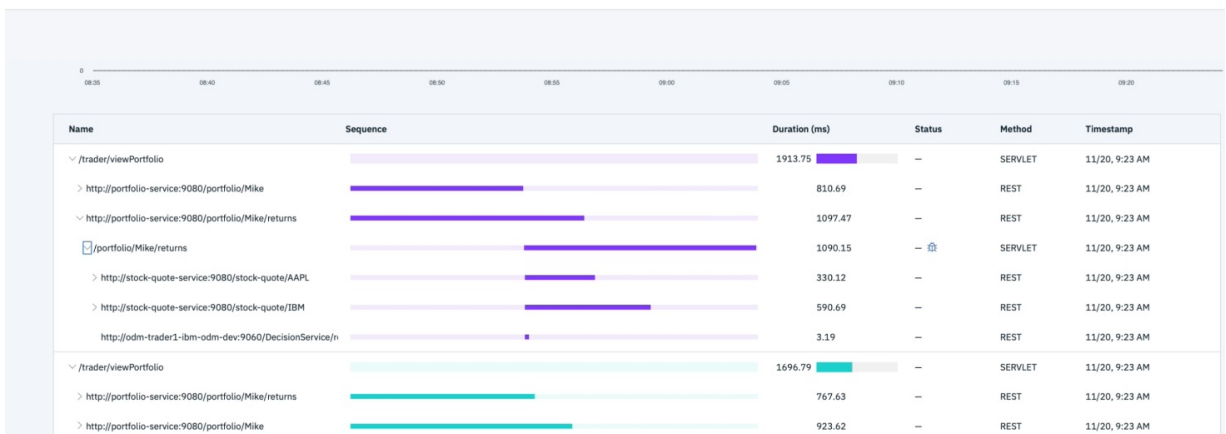


**Is this latency problem in the service itself? Compare a slow transaction to a good one to see whether the service is the problem.**

To check this hypothesis, John can choose a slow individual transaction in the transaction tracing view and compare the transaction trace with a good transaction. This is known in the industry as choosing an exemplar of a good transaction and a bad one to show the call trace history, which will point out the differences to a developer. Note: If John's access is restricted to the service and not the cluster (RBAC), he is unable to drill down to the node or cluster.

Resources / Kubernetes Service

portfolio-service



### Incident Resolution Flow for a traditional VM-based App

Resolving an incident in a traditional VM-based app in IBM Cloud App Management. In this example, Todd (IT Ops - not an SRE) receives an alert from a threshold breach. Todd has thresholds set on the application's user experience metrics such as slow synthetic transactions or slow real user transactions. Todd restores the service by creating a hypothesis and following it through to see whether he isolated the problem. If not, he creates another hypothesis. This example shows how this process is accomplished in IBM Cloud App Management. Note: This incident is more difficult to resolve since Todd does not have a

deep understanding of how the app works. He is looking at the infrastructure metrics to try to find the problem.

### **Check the incident view and determine what events were already correlated.**

To check this hypothesis, Todd goes into the incident view and sees what events were already correlated for him into a single incident by IBM Cloud App Management's event logic. Todd can now focus on a single correlated incident rather than a bigger set of apparently unrelated events. Todd checks to see whether there is a suggested runbook that is associated with the incident that can resolve this problem.

### **Check the resource view and check recent and historical trends.**

To check this hypothesis, Todd selects the resource view from the incident and checks the recent and historical trends (by using the historical slider view) to see what changes to the application's metrics happened over time. Next, Todd looks for any code deployments in the event timeline that correlates with changes in the metrics.

In the resource view, go back through the event timeline to see whether there are any anomalies detected that point to root cause.

### **Is this latency problem affecting the application caused by a network issue?**

To check this hypothesis, Todd can select a wider topology view by expanding the Service dependencies view. This expanded view includes broader elements in the system, including network elements on the topology. He looks for network events on the most problematic network area, for example, the gateway and load balancer, which connect users to the application. Examples of impacts to the application include: overloaded load balancers, slow LDAP servers, and DNS issues. If John found an issue in the broader system, he can follow the established process to resolve them with the appropriate team.

### **Is this latency problem in the application? Compare a slow transaction to a good one to see whether the problem is isolated.**

To check this hypothesis, Todd can choose a slow individual transaction in the transaction tracing view and compare the transaction trace with a good transaction. This is known in the industry as choosing an exemplar of a good transaction, and a bad one to show the call trace history, which can point out the differences to a developer. Note: If John's access is restricted to the application and not the cluster (RBAC), he is unable to drill down to the node or cluster.

## **Agents, data collectors, and plug-ins**

The ICAM Agents and ICAM Data Collectors provide monitoring in your cloud environment.

The ICAM Agents run in the on-premises application environment; ICAM Data Collectors run in the cloud environment locally or remotely. The agents and data collectors connect to the Cloud App Management server running on the IBM Cloud Private platform to monitor your business applications. You view the monitoring data and manage incidents on the Cloud App Management console. For instructions on installing agents, see [Chapter 13, "Deploying ICAM Agents,"](#) on page 193. For instructions on installing data collectors, see [Chapter 15, "Deploying ICAM Data Collectors,"](#) on page 555.

By Unified Agent, you can deploy multiple plug-ins to monitor cloud resources. Unified Agent provides a lightweight plug-in architecture, supports cloud native environment, and is easy to expand. Unified Agent can collect, process, aggregate, and write metrics to your Cloud App Management environment. It is based on Telegraf. For instructions on deploying the Unified Agent, see [Chapter 16, "Deploying Unified Agent,"](#) on page 645.

Cloud App Management provides a capability of digital experience monitoring(DEM) that can monitor web-based resources and real user experience. DEM can discover and track traffic, user behavior, and other metrics to help analyze the application performance and usability. You can enable DEM for Liberty data collector, and install the DEM plug-in for HTTP Server to monitor IBM HTTP Server and Apache HTTP Server. For more information about DEM, see [Chapter 17, "Deploying digital experience monitoring\(DEM\),"](#) on page 665.

Each agent, data collector, and plug-in monitors the resources for which it is named. For example, the IBM Integration Bus agent monitors IBM Integration Bus resources. For agent and data collector descriptions, see [“Descriptions” on page 54](#). To find out the change history of each agent and data collector, see [“Change history” on page 52](#).

The agents and data collectors for the applications that you want to monitor are available for download from Passport Advantage. There are four downloads available, one for the ICAM Agents, one for the ICAM Data Collectors, one for the Unified Agent, and one for the ICAM Extension Pack. In the following list, agents that are included in the extension pack are indicated with an asterisk \*.

Find out the list of agents, data collectors and plug-ins that are included in the packages:

### ICAM agents

- Amazon EC2 agent
- Amazon ELB agent
- Azure Compute agent
- Cassandra agent\*
- Cisco UCS agent
- Citrix VDI agent
- **2019.4.0.2** CouchDB agent\*
- DataPower® agent
- Db2 agent
- DataStage agent\*
- Hadoop agent\*
- HTTP Server agent
- IBM Integration Bus agent
- JBoss agent
- Linux OS agent
- Linux KVM agent
- MariaDB agent\*
- Microsoft Active Directory agent
- Microsoft Cluster Server agent
- Microsoft Exchange Server agent
- Microsoft .NET agent
- Microsoft Hyper-V Server agent
- Microsoft IIS agent
- Microsoft Office 365 agent\*
- Microsoft SharePoint Server agent
- Microsoft SQL Server agent
- MongoDB agent
- MySQL agent
- NetApp Storage agent
- Oracle Database agent
- PostgreSQL agent
- RabbitMQ agent\*
- SAP agent
- SAP HANA Database agent\*

- SAP NetWeaver Java Stack agent\*
- Skype for Business Server agent
- **2019.4.0.2** Sterling Connect Direct agent\*
- **2019.4.0.2** Sterling File Gateway agent\*
- Sybase agent
- Synthetics PoP
- Tomcat agent
- UNIX OS agent
- VMware VI agent
- WebLogic agent
- WebSphere Applications agent
- WebSphere Infrastructure Manager agent
- IBM MQ(formerly WebSphere MQ) agent
- Windows OS agent

#### **ICAM data collectors**

- Go data collector
- HMC agent
- J2SE data collector
- Kubernetes data collector
- Liberty data collector
- Node.js data collector
- Python data collector
- Ruby data collector

#### **Unified Agent (UA)**

- UA plug-in for Jaeger and Zipkin
- UA plug-in for NGINX
- UA plug-in for Redis
- UA plug-in for IBM API Connect(APIC)
- UA plug-in for IBM App Connect Enterprise(ACE)
- UA plug-in for IBM MQ
- UA plug-in for DEM
- UA plug-in for OpenShift

#### **DEM**

- DEM for Liberty data collector:  
It can be enabled when configuring the Liberty data collector.
- DEM for HTTP server:  
It requires 2 plug-ins:
  - UA plug-in for DEM that can be deployed by the Unified Agent package.
  - DEM plug-in for HTTP Server that can be deployed by the data collector package.

\* Extension pack agents

## Change history

Find out the information about versions and change history for each agent, data collector and plug-in.

The following table lists the agent, data collector, and plug-in names with change history technote links. Click the links to view change history details.

*Table 2. Agent , data collector, and plug-in change history*

<b>Agents, data collectors, and plug-ins</b>	<b>Links</b>
Amazon EC2 agent	<a href="#">Change history</a>
Amazon ELB agent	<a href="#">Change history</a>
Azure Compute agent	<a href="#">Change history</a>
Cassandra agent	<a href="#">Change history</a>
Cisco UCS agent	<a href="#">Change history</a>
Citrix VDI agent	<a href="#">Change history</a>
<b>2019.4.0.2</b> CouchDB agent	<a href="#">Change history</a>
DataPower agent	<a href="#">Change history</a>
DataStage agent	<a href="#">Change history</a>
Db2 agent	<a href="#">Change history</a>
DEM plug-in for HTTP Server	<a href="#">Change history</a>
Go data collector	<a href="#">Change history</a>
Hadoop agent	<a href="#">Change history</a>
HMC agent	<a href="#">Change history</a>
HTTP Server agent	<a href="#">Change history</a>
Kubernetes data collector	<a href="#">Change history</a>
IBM Integration Bus agent	<a href="#">Change history</a>
J2SE data collector	<a href="#">Change history</a>
JBoss agent	<a href="#">Change history</a>
Kubernetes data collector	<a href="#">Change history</a>
Liberty data collector	<a href="#">Change history</a>
Linux KVM agent	<a href="#">Change history</a>
Linux OS agent	<a href="#">Change history</a>
MariaDB agent	<a href="#">Change history</a>
Microsoft Active Directory agent	<a href="#">Change history</a>
Microsoft Cluster Server agent	<a href="#">Change history</a>
Microsoft Exchange Server agent	<a href="#">Change history</a>
Microsoft Hyper-V Server agent	<a href="#">Change history</a>
Microsoft IIS agent	<a href="#">Change history</a>
Microsoft .NET agent	<a href="#">Change history</a>



Table 2. Agent , data collector, and plug-in change history (continued)

<b>Agents, data collectors, and plug-ins</b>	<b>Links</b>
Microsoft Office 365 agent	<a href="#">Change history</a>
Microsoft SharePoint Server agent	<a href="#">Change history</a>
Microsoft SQL Server agent	<a href="#">Change history</a>
MongoDB agent	<a href="#">Change history</a>
MySQL agent	<a href="#">Change history</a>
NetApp Storage agent	<a href="#">Change history</a>
Node.js data collector	<a href="#">Change history</a>
Oracle Database agent	<a href="#">Change history</a>
PostgreSQL agent	<a href="#">Change history</a>
Python data collector	<a href="#">Change history</a>
RabbitMQ agent	<a href="#">Change history</a>
Ruby data collector	<a href="#">Change history</a>
SAP agent	<a href="#">Change history</a>
SAP HANA Database agent	<a href="#">Change history</a>
SAP NetWeaver Java Stack agent	<a href="#">Change history</a>
Skype for Business Server agent	<a href="#">Change history</a>
<b>2019.4.0.2</b> Sterling Connect Direct agent	<a href="#">Change history</a>
<b>2019.4.0.2</b> Sterling File Gateway agent	<a href="#">Change history</a>
Sybase agent	<a href="#">Change history</a>
Tomcat agent	<a href="#">Change history</a>
UA plug-in for DEM	<a href="#">Change history</a>
UA plug-in for IBM API Connect	<a href="#">Change history</a>
UA plug-in for IBM App Connect Enterprise	<a href="#">Change history</a>
UA plug-ins for Jaeger and Zipkin	<a href="#">Change history</a>
UA plug-in for IBM MQ	<a href="#">Change history</a>
UA plug-in for NGINX	<a href="#">Change history</a>
UA plug-in for OpenShift	<a href="#">Change history</a>
UA plug-in for Redis	<a href="#">Change history</a>
UNIX OS agent	<a href="#">Change history</a>
VMware VI agent	<a href="#">Change history</a>
WebLogic agent	<a href="#">Change history</a>
WebSphere Applications agent	<a href="#">Change history</a>
WebSphere Infrastructure Manager agent	<a href="#">Change history</a>

<i>Table 2. Agent , data collector, and plug-in change history (continued)</i>	
<b>Agents, data collectors, and plug-ins</b>	<b>Links</b>
WebSphere MQ agent	<a href="#">Change history</a>
Windows OS agent	<a href="#">Change history</a>

## Descriptions

The agent and data collector descriptions provide information about what each type of Cloud App Management agent and data collector monitors, and has links to more information.

Each agent and data collector has a version number, which changes each time the agent is updated. In any release, new agents and data collectors might be added, and existing agents might be updated. If you do not have the latest version of an agent, consider updating it. For information about how to check the version of an agent in your environment, see [Agent version command](#).

Each agent description contains information about specific agent capabilities and links to the agent configuration information.

### Amazon EC2 agent

The Monitoring Agent for Amazon EC2 offers a central point of management for your Amazon EC2 environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Amazon EC2 agent you can easily collect and analyze Amazon EC2 specific information. For more information, see [“Configuring Amazon EC2 monitoring” on page 232](#).

### Amazon ELB agent

The Monitoring Agent for AWS Elastic Load Balancer offers a central point of management for your AWS Elastic Load Balancer environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Amazon ELB agent you can easily collect and analyze AWS Elastic Load Balancer specific information. For more information, see [“Configuring AWS Elastic Load Balancer monitoring” on page 238](#).

### Azure Compute agent

The Monitoring Agent for Azure Compute offers a central point of management for your Azure Compute environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Azure Compute agent you can easily collect and analyze Azure Compute specific information.. For more information, see [“Configuring Azure Compute monitoring” on page 242](#).

### Cassandra monitoring

The Monitoring Agent for Cassandra provides you with the capability to monitor the health and performance of Cassandra cluster resources, such as nodes, keyspaces and column families. For information about configuring the agent after installation, see [“Configuring Cassandra monitoring” on page 257](#).

### Cisco UCS monitoring

The Monitoring Agent for Cisco UCS provides you with an environment to monitor the health, network, and performance of Cisco Unified Computing Systems (UCS). The Cisco UCS agent provides a comprehensive way for collecting and analyzing information that is specific to Cisco UCS and required to detect problems early and prevent them. For information about configuring the agent after installation, see [“Configuring Cisco Unified Computing System \(UCS\) monitoring” on page 250](#).

## Citrix VDI agent

The Monitoring Agent for Citrix Virtual Desktop Infrastructure monitors the following functions: Citrix XenDesktop component, Event log and alerts, and Citrix XenDesktop services. Additionally, you can view the Load Index Summary metrics performance data for Citrix XenApp and XenDesktop. You can diagnose problematic login times by viewing the performance data for the login steps. For more information, see [“Configuring Citrix Virtual Desktop Infrastructure monitoring” on page 259](#).

## CouchDB agent

**2019.4.0.2** The Monitoring Agent for CouchDB offers a central point of management for your CouchDB environment or application. The software provides a comprehensive means for gathering the information required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Monitoring Agent for CouchDB you can easily collect and analyze CouchDB specific information. For more information, see [“Configuring CouchDB monitoring” on page 266](#).

## Kubernetes data collector monitoring

With the Kubernetes data collector, you can visualize your Kubernetes environment in dashboards that show resource utilization over time to help you understand how changes to Kubernetes resources on each cluster might be impacting downstream applications. You can also monitor NGINX and Redis workloads if your environment includes them.

- For information about configuring the data collector, see [Chapter 15, “Deploying ICAM Data Collectors,” on page 555](#).
- For information about the **Resource** dashboards, see [“Viewing your managed resources” on page 769](#).

## DataPower monitoring

The Monitoring Agent for DataPower provides a central point of monitoring for the DataPower appliances in your enterprise environment. You can identify and receive notifications about common problems with the appliances. The agent also provides information about performance, resource, and workload for the appliances.

For information about configuring the agent after installation, see [“Configuring DataPower monitoring” on page 278](#).

## Db2 monitoring

The Monitoring Agent for Db2 offers a central point of monitoring for your Db2 environment. You can monitor a multitude of servers from a single IBM Cloud App Management console, with each server monitored by a Db2 agent. You can collect and analyze information in relation to applications, databases, and system resources.

For information about configuring the agent after installation, see [“Configuring Db2 monitoring” on page 282](#).

## DEM monitoring

DEM can be enabled on Liberty data collector, and can also be enabled for HTTP server to monitor web-based resources and real user experience. It can discover and track traffic, user behavior, and other metrics to help analyze the application performance and usability. For more information, see [Chapter 17, “Deploying digital experience monitoring \(DEM\),” on page 665](#).

## Hadoop monitoring

The Monitoring Agent for Hadoop provides capabilities to monitor the Hadoop cluster in your organization. You can use the agent to collect and analyze information about the Hadoop cluster, such as status of data nodes and Java™ virtual machine, memory heap and non-heap information, and information about Hadoop nodes, file systems, and queues.

For information about configuring the agent after installation, see [“Configuring Hadoop monitoring” on page 292](#).

## HMC monitoring

The Monitoring Agent for HMC provides you with the capability to monitor the Hardware Management Console (HMC). The HMC agent monitors the availability and health of HMC resources such as CPU, memory, storage, and network. It collects the following metrics: HMC, Managed Server(CEC), LPAR, VIOS, CPUPool, VSCSI, FibreChannel, and NPIV and sends these metrics to the Cloud App Management server. The supported HMC versions are V8.2 to V8.7, and V9.1. For information about configuring the agent after installation, see [Installing and Configuring the HMC agent](#).

## HTTP Server monitoring

The Monitoring Agent for HTTP Server collects performance data about the IBM HTTP Server. For example, server information, such as the status and type of server, the number of server errors, and the number of successful and failed logins to the server are shown. A data collector gathers the data that is sent to the HTTP Server agent. The agent runs on the same system with the IBM HTTP Server that it monitors. Each monitored server is registered as a subnode. For more information, see [“Configuring HTTP Server agent monitoring” on page 301](#).

## IBM Integration Bus monitoring

The Monitoring Agent for IBM Integration Bus is a monitoring and management tool that provides you with the means to verify, analyze, and tune message broker topologies that are associated with the IBM WebSphere Message Broker and IBM Integration Bus products.

For information about configuring the agent after installation, see [“Configuring IBM Integration Bus monitoring” on page 304](#).

## InfoSphere DataStage monitoring

The Monitoring Agent for InfoSphere DataStage offers a central point of management for InfoSphere DataStage application's Service Tier as well as Engine Tier. You can use the InfoSphere DataStage agent to monitor details, such as Job Runs, CPU and Memory of engines, historical trend, status of services, and so on. Information is standardized across the system. You can monitor multiple engines from a single point. By using the InfoSphere DataStage Application agent you can easily collect and analyze InfoSphere DataStage Application specific information.

For information about configuring the agent after installation, see [“Configuring InfoSphere DataStage monitoring” on page 313](#).

## JBoss agent

The Monitoring Agent for JBoss offers a central point of management for your JBoss environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the JBoss agent you can easily collect and analyze JBoss specific information. For more information, see [“Configuring JBoss monitoring” on page 316](#).

## Go monitoring

The Go data collector can provide you with visibility and control of your Go applications, and help you ensure optimal performance and efficient use of resources. You can reduce and prevent application crashes and slowdowns around the clock, as the data collector assists you in detecting, diagnosing and isolating performance issues. For more information, see [“Configuring Go application monitoring” on page 579](#).

## J2SE monitoring

The J2SE data collector is a greenfield runtime data collector that monitors the cloud-based Java applications. The J2SE data collector helps you to manage the performance and availability of stand-alone Java applications in IBM Cloud Private.

For information about configuring the data collector, see [“Configuring J2SE application monitoring” on page 584](#).

## Liberty monitoring

The Liberty data collector monitors the Liberty applications or Microclimate-based Liberty applications in IBM Cloud Private.

### IBM Cloud Private applications

- For more information about configuring the data collector in IBM Cloud Private and Microclimate, see [Monitoring Liberty applications in IBM Cloud Private](#) and [Monitoring Microclimate-based Liberty applications in IBM Cloud Private](#).

## Linux KVM monitoring

The Monitoring Agent for Linux KVM offers a central point of management for your Linux Kernel-based Virtual Machines environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single IBM Cloud App Management console. By using the Linux KVM agent you can easily collect and analyze Linux Kernel-based Virtual Machines specific information.

For information about configuring the agent after installation, see [“Configuring Linux KVM monitoring” on page 324](#).

## Linux OS monitoring

The Monitoring Agent for Linux OS provides monitoring capabilities for the availability, performance, and resource usage of the Linux OS environment. This agent supports Docker container monitoring. For example, detailed information such as the CPU usage, memory, network and I/O usage information that relates to the docker container is shown. General information about the docker containers running on the server, such as the docker ID and instance name is also shown. You can collect and analyze server-specific information, such as operating system and CPU performance, Linux disk information and performance analysis, process status analysis, and network performance.

## MariaDB monitoring

The Monitoring Agent for MariaDB offers a central point of management for your MariaDB environment or application. The software provides a comprehensive means for gathering the information required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Monitoring Agent for MariaDB you can easily collect and analyze MariaDB specific information.

For information about configuring the agent after installation, see [“Configuring MariaDB monitoring” on page 335](#)

## Microsoft Active Directory monitoring

The Monitoring Agent for Microsoft Active Directory provides capabilities to monitor the Active Directory in your organization. You can use the agent to collect and analyze information that is specific to Active Directory.

For information about configuring the agent after installation, see [“Configuring Microsoft Active Directory monitoring” on page 340](#)

## Microsoft Cluster Server monitoring

The Monitoring Agent for Microsoft Cluster Server provides capabilities to monitor the Microsoft Cluster Server in your organization. You can use the Microsoft Cluster Server agent to collect information that is related to cluster resource availability, such as cluster level, cluster nodes, cluster resource groups, cluster resources, and cluster networks. The agent also provides statistics for cluster resources usage, such as processor usage, memory usage, disk usage, and network usage.

For information about configuring the agent after installation, see [“Configuring Microsoft Cluster Server monitoring” on page 344](#)

## Microsoft .NET monitoring

The Monitoring Agent for Microsoft .NET offers a central point of management for your Microsoft .NET environment or application. With the Monitoring Agent for Microsoft .NET, you can easily collect and

analyze Microsoft .NET specific information from Cloud App Management console. The agent also monitors various applications, services and processes that uses the .Net CLR. For information about configuring the agent after installation, see [“Configuring Microsoft .NET monitoring”](#) on page 338.

### **Microsoft Exchange Server monitoring**

The Monitoring Agent for Microsoft Exchange Server provides capabilities to monitor the health, availability, and performance of the Exchange Servers in your organization. You can use the Microsoft Exchange Server agent to collect server-specific information, such as mail traffic, state of mailbox databases and activities of clients. Additionally, the agent provides statistics of cache usage, mail usage, database usage and client activities to help you analyze the performance of Exchange Servers. For information about configuring the agent after installation, see [“Configuring Microsoft Exchange Server monitoring”](#) on page 346.

### **Microsoft Hyper-V Server monitoring**

The Monitoring Agent for Microsoft Hyper-V Server provides capability to monitor the availability and performance of all the Hyper-V systems in your organization. The Microsoft Hyper-V Server agent provides configuration information such as the number of virtual machines, the state of the virtual machines, the number of allocated virtual disks, the allocated virtual memory, and so on. Additionally, the agent provides statistics of physical processor usage, memory usage, network usage, logical processor usage, and virtual processor usage. For information about configuring the agent after installation, see [“Configuring Microsoft Hyper-V monitoring”](#) on page 358.

### **Microsoft IIS monitoring**

The Monitoring Agent for Microsoft Internet Information Services offers a central point of management for your Microsoft Internet Information Server environment or application. You can use the Microsoft Internet Information Server agent to monitor website details such as request rate, data transfer rate, error statistics, and connections statistics. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Microsoft Internet Information Server agent you can easily collect and analyze Microsoft Internet Information Server specific information. For information about configuring the agent after installation, see [“Configuring Microsoft IIS monitoring”](#) on page 362.

### **Microsoft Office 365 monitoring**

The Monitoring Agent for Microsoft Office 365 provides capabilities to monitor your Microsoft Office 365 environment or application. You can use the Microsoft Office 365 agent to monitor the health and performance of Office 365 resources, such as the Office 365 subscribed services, Office 365 portal, mailbox users, SharePoint sites, and OneDrive storage. For information about configuring the agent after installation, see [“Configuring Microsoft Office 365 monitoring”](#) on page 364.

### **Microsoft SharePoint Server monitoring**

The Monitoring Agent for Microsoft SharePoint Server provides you with the environment to monitor the availability, events, and performance of the Microsoft SharePoint Server. Use this agent to gather data from the Microsoft SharePoint Server and manage operations. For information about configuring the agent after installation, see [“Configuring Microsoft SharePoint Server monitoring”](#) on page 370.

### **Microsoft SQL Server monitoring**

The Monitoring Agent for Microsoft SQL Server offers a central point of monitoring for your Microsoft SQL Server environment or application. You can collect and analyze Microsoft SQL Server specific information, and monitor multiple servers from a single IBM Cloud App Management console. For information about configuring the agent after installation, see [“Configuring Microsoft SQL Server monitoring”](#) on page 373.

## MongoDB monitoring

The Monitoring Agent for MongoDB provides monitoring capabilities for the usage, status, and performance of the MongoDB deployment. You can collect and analyze information such as database capacity usage, percentage of connections open, memory usage, instance status, and response time in visualized dashboards.

For information about configuring the agent after installation, see [“Configuring MongoDB monitoring” on page 401](#).

## MySQL monitoring

The Monitoring Agent for MySQL provides monitoring capabilities for the status, usage, and performance of the MySQL deployment based on the 2 top level resources classification.

- MySQL Database

On the Database Resource, you can view and analyze the information specific to different databases of your MySQL Server, such as Table Count, Database Size, ProcessList Details, Events data and others.

- MySQL Instance

On the Resource page, you can view all the instances for MySQL and analyze information such as Percentage of Active Connections, Slow Queries, Bytes Received vs Sent, Error Details, CPU, Memory Usage and others.

For information about configuring the agent after installation, see [“Configuring MySQL monitoring” on page 406](#).

## NetApp Storage monitoring

The Monitoring Agent for NetApp Storage provides capabilities to monitor your NetApp storage systems by using the NetApp OnCommand Unified Manager (OCUM). You can use the NetApp Storage agent to monitor the health and performance of ONTAP cluster with event-driven responses and precise representation of historical trends in the Cloud App Management console.

For information about configuring the agent after installation, see [“Configuring NetApp Storage monitoring” on page 408](#).

## Node.js monitoring

The Node.js data collector monitors the Node . js applications or Microclimate-based Node . js applications in IBM Cloud Private. For more information, see [“Configuring Node.js application monitoring” on page 601](#).

## Oracle Database monitoring

The Monitoring Agent for Oracle Database provides monitoring capabilities for the availability, performance, and resource usage of the Oracle database. You can configure more than one Oracle Database agent instance to monitor different Oracle databases. Remote monitoring capability is also provided by this agent. For more information see, [“Configuring Oracle Database monitoring” on page 414](#).

## PostgreSQL monitoring

The Monitoring Agent for PostgreSQL monitors the PostgreSQL database by collecting PostgreSQL metrics through a JDBC driver. The agent provides data about system resource usage, database capacity, connections that are used, individual status of running instances, statistics for operations, response time for SQL query statements, database size details, and lock information.

For information about configuring the agent after installation, see [“Configuring PostgreSQL monitoring” on page 450](#).

## Python monitoring

The Python data collector is a runtime data collector that helps you monitor Python applications to ensure optimal performance and efficient use of resources, reduce, and prevent application crashes and slowdowns in IBM Cloud Private.

For information about configuring the data collector, see [“Configuring Python application monitoring” on page 606](#).

### **RabbitMQ monitoring**

The Monitoring Agent for RabbitMQ provides you with the capability to monitor the RabbitMQ cluster. You can collect and analyze information about the nodes, queues, and channels of the RabbitMQ cluster.

For information about configuring the agent after installation, see [“Configuring RabbitMQ monitoring” on page 454](#).

### **Ruby monitoring**

The Ruby data collector can provide you with visibility and control of the Ruby application, and help you ensure optimal performance and efficient use of resources. For more information, see [“Configuring Ruby application monitoring” on page 617](#).

### **Synthetics PoP**

Use the Synthetics PoP to monitor your REST calls and other urls from multiple locations. Create synthetic tests and schedule them to run on a predefined schedule. Monitor both the availability and response time of you websites. For more information, see [“Synthetics PoP” on page 620](#).

### **SAP monitoring**

The Monitoring Agent for SAP Applications provides the capability to monitor your SAP system. The SAP agent offers a central point of management for gathering the information to detect problems early, and prevent them. It enables effective systems management across SAP releases, applications, components, and the underlying databases, operating systems, and external interfaces.

For more information about configuring the agent after installation, see [“Configuring SAP monitoring” on page 456](#).

### **SAP HANA Database monitoring**

The Monitoring Agent for SAP HANA Database monitors availability, resource usage, and performance of the SAP HANA database. It can monitor HANA deployment scenarios such as single host - single database, single host - multiple tenant databases, multiple hosts - single database, and multiple hosts - multiple tenant databases. You can analyze the information that the agent collects and take appropriate actions to resolve issues in the SAP HANA Database.

For more information about configuring the agent after installation, see [“Configuring SAP HANA Database monitoring” on page 485](#).

### **SAP NetWeaver Java Stack monitoring**

The Monitoring Agent for SAP NetWeaver Java Stack monitors the availability, resource usage, and performance of the SAP NetWeaver Java Stack. The agent can monitor SAP NetWeaver Java Stack deployment scenarios such as single host - single instance, single host - multiple instances, multiple hosts - single instances, and multiple hosts - multiple instances. You can analyze the information that the agent collects and take appropriate actions to resolve issues in the SAP NetWeaver Java Stack.

For information about configuring the agent after installation, see [“Configuring SAP NetWeaver Java Stack monitoring” on page 488](#).

### **Skype for Business Server monitoring**

The Monitoring Agent for Skype for Business Server provides you with the capability to monitor the Skype for Business Server. You can use the agent to monitor the availability, performance, error log, event log, and historical data of the Business Server.

For information about configuring the agent after installation, see [“Configuring Skype for Business Server monitoring” on page 491](#).

### **Sybase Server monitoring**

The Sybase agent offers a central point of management for distributed databases. It collects the required information for database and system administrators to examine the performance of the Sybase server system, detect problems early and prevent them.



For information about configuring the agent after installation, see [“Configuring Sybase Server monitoring”](#) on page 505.

### **Sterling Connect Direct monitoring**

**2019.4.0.2** The Monitoring Agent for Sterling Connect Direct monitors the health and performance of Connect Direct nodes in your organization. By using the Sterling Connect Direct agent, you can easily analyze the file transfer activity within and between organizations. For information about configuring the agent after installation, see [“Configuring Sterling Connect Direct monitoring”](#) on page 496.

### **Sterling File Gateway monitoring**

**2019.4.0.2** The Monitoring Agent for Sterling File Gateway monitors the Sterling File Gateway application, which is used for transferring files between internal and external partners by using different protocols, different file naming conventions, and different file formats. It also supports the remote monitoring feature. For information about configuring the agent after installation, see [“Configuring Sterling File Gateway monitoring”](#) on page 499.

### **Tomcat monitoring**

The Monitoring Agent for Tomcat offers a central point of management for your Tomcat environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single IBM Cloud App Management console. By using the Tomcat agent you can easily collect and analyze Tomcat specific information. For information about configuring the agent after installation, see [“Configuring Tomcat Monitoring”](#) on page 511

### **Unified Agent monitoring**

By Unified Agent, you can deploy multiple plug-ins to monitor cloud resources. Unified Agent provides a lightweight plug-in architecture, supports cloud native environment, and is easy to expand. Unified Agent can collect, process, aggregate, and write metrics to your Cloud App Management environment. It is based on Telegraf. For instructions on deploying the Unified Agent, see [Chapter 16, “Deploying Unified Agent,”](#) on page 645.

### **UNIX OS monitoring**

The Monitoring Agent for UNIX OS provides monitoring capabilities for the availability, performance, and resource usage of the UNIX OS environment. (AIX operating system only. See [“System requirements”](#) on page 75.) You can collect and analyze server-specific information, such as operating system and CPU performance, UNIX disk information and performance analysis, process status analysis, and network performance.

### **VMware VI monitoring**

The Monitoring Agent for VMware VI offers a central point of management for your VMware VI environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single IBM Cloud App Management console. By using the VMware VI agent you can easily collect and analyze VMware specific information. For information about configuring the agent after installation, see [“Configuring VMware VI monitoring”](#) on page 515.

### **WebLogic monitoring**

The Monitoring Agent for WebLogic provides you with a central point of monitoring for the health, availability, and performance of your WebLogic server environment. The agent displays a comprehensive set of metrics to help you make informed decisions about your WebLogic resources, including Java virtual machines (JVMs), Java messaging service (JMS), Java Database Connectivity

(JDBC). For information about configuring the agent after installation, see [“Configuring WebLogic monitoring”](#) on page 270.

### **WebSphere Applications monitoring**

The Monitoring Agent for WebSphere Applications with the embedded data collector monitors the resources of WebSphere application servers. These monitoring components can be configured to do the following things:

- Gather PMI metrics for resource monitoring through a JMX interface on the application server.
- Gather aggregated request performance metrics.
- Track the performance of individual request and method calls.

The monitoring data is displayed in the Cloud App Management user interface. You can use the provided dashboards to isolate specific problem areas of your application server. Drill down to determine whether a problem lies with an underlying resource or if it relates to the application's code.

For information about configuring the agent after installation, see [“Configuring WebSphere Applications monitoring”](#) on page 523.

### **WebSphere Infrastructure Manager monitoring**

The Monitoring Agent for WebSphere Infrastructure Manager provides the monitoring capabilities for the WebSphere Application Server Deployment Manager and Node Agent, including server status, resources, and transactions. You can use the data that is collected by the WebSphere Infrastructure Manager agent to analyze the performance of your Deployment Manager and Node Agent, and whether a problem occurred.

For information about configuring the agent after installation, see [“Configuring WebSphere Infrastructure Manager monitoring”](#) on page 547.

### **WebSphere MQ monitoring**

With the Monitoring Agent for IBM MQ(formerly IBM WebSphere® MQ), you can easily collect and analyze data that is specific to WebSphere MQ for your queue managers from a single vantage point. You can then track trends in the data that is collected and troubleshoot system problems by using the predefined dashboards.

For information about configuring the agent after installation, see [“Configuring WebSphere MQ monitoring”](#) on page 548.

### **Windows OS monitoring**

The Monitoring Agent for Windows OS provides monitoring capabilities for the availability, performance, and resource usage of the Windows OS environment. You can collect and analyze server-specific information, such as operating system and CPU performance, disk information and performance analysis, process status analysis, Internet session data, monitored logs information, Internet server statistics, message queuing statistics, printer and job status data, Remote Access Services statistics, and services information.

---

## Chapter 5. Accessibility features

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

### Accessibility features

The web-based interface of IBM Cloud App Management is the Cloud App Management console. The console includes the following major accessibility features:

- Enables users to use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen. Consult the product documentation of the assistive technology for details on using those technologies with this product.
- Enables users to operate specific or equivalent features using only the keyboard.
- Communicates all information independently of color.<sup>1</sup>

The Cloud App Management console uses the latest W3C Standard, WAI-ARIA 1.0 (<http://www.w3.org/TR/wai-aria/>), to ensure compliance with US Section 508 (<http://www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards>), and Web Content Accessibility Guidelines (WCAG) 2.0 (<http://www.w3.org/TR/WCAG20/>). To take advantage of accessibility features, use the latest release of your screen reader in combination with the latest web browser that is supported by this product.

The Cloud App Management console online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described at [IBM Knowledge Center release notes](http://www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility) [http://www.ibm.com/support/knowledgecenter/doc/kc\\_help.html#accessibility](http://www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

### Keyboard navigation

This product uses standard navigation keys.

### Interface information

The Cloud App Management console web user interface does not rely on cascading style sheets to render content properly and to provide a usable experience. However, the product documentation does rely on cascading style sheets. IBM Knowledge Center provides an equivalent way for low-vision users to use their custom display settings, including high-contrast mode. You can control font size by using the device or browser settings.

The Cloud App Management console web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

The Cloud App Management console user interface does not have content that flashes 2 - 55 times per second.

### Related accessibility information

In addition to standard IBM help desk and support websites, IBM has established a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service 800-IBM-3383 (800-426-3383) (within North America)

### IBM and accessibility

For more information about the commitment that IBM has to accessibility, see [IBM Accessibility](http://www.ibm.com/able) ([www.ibm.com/able](http://www.ibm.com/able)).

---

<sup>1</sup> Exceptions include some **Agent Configuration** pages of the Performance Management console.



---

## Chapter 6. Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. 2018.

## Trademarks

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

## Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

### Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

### Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

### Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

### Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

## IBM Online Privacy Statement

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth in the following paragraphs.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's user name for purposes of session management, authentication, and single sign-on configuration. These cookies can be disabled, but disabling them will also likely eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.



---

# Chapter 7. Planning your deployment

To ensure that your IBM Cloud App Management deployment is successful, planning is critical. A successful deployment can be completed in a few main steps. Ensure you complete all the procedures in each step. This planning deployment scenario assumes that you are an administrator working on a Linux 64-bit system.

Complete the following steps:

- [“Step 1: Determine your hardware requirements” on page 69](#)
- [“Step 2: Determine the storage type to use” on page 69](#)
- [“Step 3: Deploy IBM Cloud Private Enterprise V3.2.1” on page 70](#)
- [Step 4: Download and deploy the Cloud App Management server](#)
- [“Step 5: Access the Cloud App Management console” on page 70](#)
- [Step 6: Deploy the agents](#)

## Step 1: Determine your hardware requirements

### Do you plan to install <100 monitored resources?

Use a small demonstration (trial or proof of concept) environment. You will need the following resources: 1 VMs, 8 CPU, 32 GB memory, Disk space 100 GB.

### Do you plan to install >100 monitored resources?

The required hardware resources will vary depending on metrics per minute. Determine the metrics per minute value:

1. In the estimation spreadsheet [IBM Cloud App Management Load Projections Spreadsheet](#), enter an estimate of the number of monitored resources you plan to install; a metrics per minute value is returned.
2. Based on the returned metrics per minute value, review the [“Planning hardware and sizing” on page 77](#) topic to determine the required number of VMs, CPU, memory and disk space.

## Step 2: Determine the storage type to use

You will use the `prepare-pv.sh` script to create the storage class and PVs for your statefulset services. You must create the PVs for the Cassandra, Kafka, ZooKeeper, CouchDB, and Datalayer statefulset services in your deployment. During the Cloud App Management server install, when you run the `prepare-pv.sh` script, you must select your storage type either local or vSphere.

### Local storage

Local storage uses local persistent volume storage. Local storage is recommended. Local storage is similar to `hostpath`, in that it creates a PersistentVolume (PV) that uses a local directory on a system. Local storage differs from `hostpath` because of its affinity to lock storage on the node and the local directory. This affinity prevents the statefulset pod being moved to another system and losing its storage. With `hostPath`, if the pod is moved, you can lose persistence. Local storage offers better performance than network-based storage like NFS or GlusterFS. With local storage, if the node or its storage gets lost, then the PV data is also lost.

For local storage, the IP is the address of the IBM Cloud Private worker node that you want to assign to a specific PV. The PV and any service, which is claiming that particular node, is permanently locked to that node. If you lose that worker node, you lose that PV and the service. If your IBM Cloud Private cluster is running on a VMware deployment with administrator access to vSphere storage, VMware drives are automatically provisioned and locally attached to the correct VM. If the VM fails, the drive is moved to a new VM with the statefulset pod. For more information about setting up vSphere storage in IBM Cloud Private, see the [vSphere Cloud Provider](#) topic.

## vSphere

vSphere storage uses vSphere provisioned storage. vSphere requires existing vSphere storage class. If your IBM Cloud Private cluster is running on a VMware deployment with administrator access to vSphere storage, VMware drives are automatically provisioned and locally attached to the correct VM. If the VM fails, the drive is moved to a new VM with the statefulset pod. For more information about setting up vSphere storage in IBM Cloud Private, see the [vSphere Cloud Provider](#) topic.

### Step 3: Deploy IBM Cloud Private Enterprise V3.2.1

For the preparation steps that you must complete before you install IBM Cloud Private Enterprise and the instructions to download and install IBM Cloud Private Enterprise, see the [Chapter 8, “Deploying IBM Cloud Private,”](#) on page 89 section.

### Step 4: Download and deploy the Cloud App Management server

Complete the procedures in the [“Offline: Installing IBM Cloud App Management stand-alone on IBM Cloud Private”](#) on page 148 section.

### Step 5: Access the Cloud App Management console

Complete the steps in the [“Starting the Cloud App Management UI”](#) on page 176 topic.

### Step 6: Deploy the agents and data collectors

#### Do you have IBM Tivoli Monitoring agents or IBM Tivoli Composite Application Manager agents (referred to as V6 agents) connecting to the Tivoli Enterprise Monitoring Server?

Configure these agents to connect to the Cloud App Management server. You can then view monitoring data on the Cloud App Management console. Complete the procedures in the [“Integrating with IBM Tivoli Monitoring agents”](#) on page 675 section.

#### Do you have Cloud APM V8.1.4 agents (referred to as V8 agents) connecting to the Cloud APM server?

Configure these agents to connect to the Cloud App Management server. You can then view monitoring data on the Cloud App Management console. Complete the procedures in the [“Integrating with Cloud APM, Private agents”](#) on page 688 section.

#### Do you have an environment with no previous V6 or V8 agents installed?

Complete the procedures in the [Chapter 13, “Deploying ICAM Agents,”](#) on page 193 topics.

For configuring ICAM Data Collectors, complete the procedures in the [Chapter 15, “Deploying ICAM Data Collectors,”](#) on page 555 topics.

## Product components

For a Cloud App Management deployment, you must download a number of components.

Before you deploy a Cloud App Management environment, you must:

- Download the installation files for IBM Cloud Private. For more information, see the [Set up the installation environment](#) section of the [Installing an IBM Cloud Private Enterprise environment](#) topic in the IBM Cloud Private Knowledge Center.
- Download the Cloud App Management installation packages from the [IBM Passport Advantage](#) website. For more information, see [“Part numbers”](#) on page 71.
- If you plan to integrate with IBM Tivoli Monitoring agents, you must download the 6.3.0.7-TIV-ITM\_TEMA-IF0008 agent patch from IBM Fix Central, for more information, see [“Connecting IBM Tivoli Monitoring agents to Cloud App Management server”](#) on page 676.
- Download the agents installation images from [IBM Passport Advantage](#). For more information, see [“Part numbers”](#) on page 71.

Other available bundled products include:

- IBM Tivoli Monitoring

**Note:** IBM Operations Analytics Log Analysis V1.3.5 is available in this bundle.

- IBM Cloud Application Performance Management, Private

**Note:** Db2 V11.1 is available in this bundle.

For details of part numbers for all bundled products, and links to the relevant knowledge centers, see [“Part numbers” on page 71](#).

## Part numbers

Review the part numbers to identify the components to download from the [IBM Passport Advantage](#) website for your IBM Cloud App Management V2019.4.0 installation.

### IBM Cloud App Management components

eImage descriptions	IBM Cloud App Management (CJ6KVEN)
IBM Cloud App Management V2019.4.0 Server Install on AMD64	CC4KNEN icam_ppa_2019.4.0_prod.tar.gz
IBM Cloud App Management V2019.4.0.2 Agents Install xLinux	CC5GXEN appMgtAgents_xlinux_2019.4.0.2.tar.gz
IBM Cloud App Management V2019.4.0.2 Agents Install zLinux	CC5H1EN appMgtAgents_zlinux_2019.4.0.2.tar.gz
IBM Cloud App Management V2019.4.0.2 Agents Install Windows	CC5GYEN appMgtAgents_win_2019.4.0.2.zip
IBM Cloud App Management V2019.4.0.2 Agents Install AIX	CC5GZEN appMgtAgents_aix_2019.4.0.2.tar.gz
IBM Cloud App Management V2019.4.0.2 Agents Install Solaris Sparc	CC5H2EN appMgtAgents_solaris_2019.4.0.2.tar.gz
IBM Cloud App Management V2019.4.0.2 Agents Install Solaris x86	CC5H5EN appMgtAgents_solaris_x862019.4.0.2.tar.gz
IBM Cloud App Management V2019.4.0.2 Agents Install PLinux	CC5H3EN appMgtAgents_plinux_2019.4.0.2.tar.gz
IBM Cloud App Management V2019.4.0.2 Agents Install PLinuxLE	CC5H4EN appMgtAgents_plinuxle_2019.4.0.2.tar.gz
IBM Cloud App Management V2019.4.0.2 Data Collectors Install	CC5H0EN appMgtDataCollectors_2019.4.0.2.tar.gz (Contains sub packages, described in <a href="#">“Data collectors sub-packages” on page 72</a> .)
Multicluster Event Management Klusterlet on PlinuxLE	CC4KXEN agent_ppa_2019.4.0_prod_ppc64le.tar.gz
Multicluster Event Management Server on PlinuxLE	CC4KYEN icam_ppa_2019.4.0_prod_lite_ppc64.tar.gz
Multicluster Event Management Klusterlet on AMD64	CC4KZEN agent_ppa_2019.4.0_prod_amd64.tar.gz
Multicluster Event Management Server on AMD64	CC4LOEN icam_ppa_2019.4.0_prod_lite_amd64.tar.gz

Table 3. Cloud App Management 2019.4.0 Multiplatform eAssembly part numbers (for both base and advanced offerings) and component part numbers (continued)

eImage descriptions	IBM Cloud App Management (CJ6KVEN)
Multicluster Event Management Klusterlet on Zlinux	CC4TGEN agent_ppa_2019.4.0_prod_s390x.tar.gz
IBM Cloud Unified Agent V2019.4.0.1	CC501EN unifiedAgent_2019.4.0.1.tar.gz

#### Data collectors sub-packages

appMgtDataCollectors\_2019.4.0.2.tar.gz contains the following 3 sub-packages:

- app\_mgmt\_k8sdc.tar.gz (contains sub-packages)
- app\_mgmt\_runtime\_dc\_2019.4.0.2.tar.gz
- app\_mgmt\_syntheticpop\_xlinux.tar.gz

app\_mgmt\_k8sdc.tar.gz contains the following sub-packages and files:

- app\_mgmt\_k8sdc\_docker.tar.gz
- app\_mgmt\_k8sdc\_operator.tar.gz
- app\_mgmt\_k8sdc\_helm
- helm-main.yaml
- README.md

#### Note:

- If you plan to integrate with IBM Tivoli Monitoring agents, you must download the 6.3.0.7-TIV-ITM\_TEMA-IF0008 agent patch from IBM Fix Central, for more information, see [“Connecting IBM Tivoli Monitoring agents to Cloud App Management server”](#) on page 676.

#### Extension packs available with Cloud App Management base and advanced

Table 4. Extension Pack file names, eAssembly part numbers (in parentheses), and eImage part numbers

eImage descriptions	IBM Cloud App Management V2019.4.0 Extension Pack (CJ6KWEN)
IBM Cloud App Management V2019.4.0.2 Extension Pack xLinux	CC5H6EN appMgtExt_xlinux_2019.4.0.2.tar
IBM Cloud App Management V2019.4.0.2 Extension Pack Windows	CC5H7EN appMgtExt_win_2019.4.0.2.zip
IBM Cloud App Management V2019.4.0.2 Extension Pack AIX	CC5H8EN appMgtExt_aix_2019.4.0.2.tar
IBM Cloud App Management V2019.4.0.2 Extension Pack PLinux	CC5H9EN appMgtExt_plinux_2019.4.0.2.tar
IBM Cloud App Management V2019.4.0.2 Extension Pack PLinuxLE	CC5HAEN appMgtExt_plinuxle_2019.4.0.2.tar

#### IBM Cloud Private components

Before you install the IBM Cloud App Management product, you must install the IBM Cloud Private platform. The part numbers and file names for the IBM Cloud Private components are in [Table 5](#) on page 73.

Table 5. IBM Cloud Private component eAssembly part numbers (in parentheses) and eImage part numbers

eImage descriptions	IBM Cloud Private (CJ5NFEN)	Product documentation
IBM Cloud Private Foundation 3.2.1 Quick Start Guide	CC3KNML	
IBM Cloud Private 3.2.1 for Linux (x86_64) Docker	CC3KPEN	<a href="#">Click here</a>
IBM Cloud Private 3.2.1 Docker for Linux (x86_64)	CC3KUEN	<a href="#">Click here</a>
IBM Cloud Private for Red Hat Enterprise Linux OpenShift (64-bit) Docker	CC3KREN	<a href="#">Click here</a>

### IBM Tivoli Monitoring bundled products

After you purchase your IBM Cloud App Management license, you can download any or all of the IBM Tivoli Monitoring bundled software. The eAssembly part numbers for the Tivoli Monitoring bundled products that are available with your offering are listed in [Table 6 on page 73](#).

Table 6. Software eAssembly part numbers for IBM Tivoli Monitoring bundled products that are downloadable from Passport Advantage for use with IBM Cloud App Management, Base and IBM Cloud App Management, Advanced

Title on Passport Advantage	eAssembly number	Product documentation
IBM Tivoli Monitoring V6.3.0.2 Quick Start Guide	CJ403ML	<a href="#">Click here</a>
IBM Tivoli Monitoring V6.3.0.7 for IBM Cloud App Management	CJ404ML	<a href="#">Click here</a>
IBM Tivoli Monitoring V6.3.0.7 Agent for IBM Cloud App Management	CJ405ML	<a href="#">Click here</a>
IBM Db2 Advanced Workgroup Server Edition 11.1 for IBM Tivoli Monitoring V6.3 for IBM Cloud App Management	CJ40AML	<a href="#">Click here</a>
IBM Tivoli Monitoring V6.3 IBM Tivoli System Automation For Multiplatform Base V3.2.0 for IBM Cloud App Management	CJ40BML	<a href="#">Click here</a>
IBM Tivoli Monitoring Version 6.3: Netcool System Service Monitor Component V4.0.1 for IBM Cloud App Management	CJ40CEN	<a href="#">Click here</a>
Jazz for Service Management V1.1.2.0 for IBM Tivoli Monitoring V6.3 for IBM Cloud App Management	CJ40DML	<a href="#">Click here</a>

### IBM Cloud Application Performance Management, Private bundled components

After you purchase your IBM Cloud App Management license, you can download any or all of the IBM Cloud Application Performance Management, Private bundled software. The eAssembly part numbers for the Cloud APM, Private bundled products that are available with your offering are listed in [Table 7 on page 74](#) (IBM Cloud App Management, Advanced).

*Table 7. Software eAssembly part numbers for Cloud APM, Private bundled products that are downloadable from Passport Advantage for use with IBM Cloud App Management, Base and IBM Cloud App Management, Advanced*

<b>Title on Passport Advantage</b>	<b>eAssembly number</b>	<b>Product documentation</b>
IBM Cloud Application Performance Management, Base Private V8.1.4 for IBM Cloud App Management Multiplatform Multilingual eAssembly	CJ40EML	<a href="#">Click here</a>
IBM Cloud Application Performance Management, Advanced Extension Pack V8.1.4 Multiplatform Multilingual eAssembly	CJ6PUEN	<a href="#">Click here</a>
IBM Cloud Application Performance Management, Base Extension Pack V8.1.4 Multiplatform Multilingual eAssembly for ICAM Cloud App Management	CJ6PVEN	<a href="#">Click here</a>
IBM Cloud Application Performance Management, Infrastructure Extension Pack V8.1.4 Multiplatform Multilingual eAssembly for ICAM Cloud App Management	CJ6PWEN	<a href="#">Click here</a>
IBM Operations Analytics Log Analysis for IBM Cloud Application Performance Management, Base Private for IBM Cloud App Management Multiplatform	CJ40FML	<a href="#">Click here</a>
IBM SmartCloud Application Performance Management Entry Edition for IBM Cloud Application Performance Management, Base Private Multiplatform for IBM	CJ40GML	<a href="#">Click here</a>
IBM Cloud Application Performance Management, Advanced Private V8.1.4 Multiplatform Multilingual eAssembly for IBM Cloud App Management	CJ4JKEN	<a href="#">Click here</a>
IBM SmartCloud Application Performance Management Standard & Non-Prod V7.7 for IBM Cloud Application Performance Management, Advanced Private	CJ4JLEN	<a href="#">Click here</a>
IBM Tivoli Composite Application Manager Transactions V7.4.0.1 Response Time and Internet Service Monitoring for IBM Cloud Application Performance	CJ4JMEN	<a href="#">Click here</a>
IBM Tivoli Composite Application Manager Transactions V7.4.0.1 Transaction Tracking for IBM Cloud Application Performance Management, Advanced Private	CJ4JNEN	<a href="#">Click here</a>
IBM Rational 8.6 for IBM Tivoli Composite Application Manager Transactions V7.4 for IBM Cloud Application Performance Management, Advanced	CJ4JPEN	<a href="#">Click here</a>
IBM Tivoli Composite Application Manager for Microsoft Applications Version 6.3.1 Advance eAssembly (ITCAMMA 6.3.1) Multiplatform, Multilingual eAssembly	CJ4JQEN	<a href="#">Click here</a>
IBM Tivoli Composite Application Manager for Applications V7.2.1.2 for IBM Cloud App Management	CJ4JREN	<a href="#">Click here</a>
IBM Tivoli Monitoring for Virtual Environments V7.2.0.3 Quick Start Guide Multilingual eAssembly for IBM Cloud App Management	CJ4JSEN	<a href="#">Click here</a>
IBM Tivoli Composite Application Manager for Application Diagnostics V7.1.0.4: Managing Server Component eAssembly, Multiplatform, Multilingual for IBM Cloud App Management	CJ4JTEN	<a href="#">Click here</a>

Table 7. Software eAssembly part numbers for Cloud APM, Private bundled products that are downloadable from Passport Advantage for use with IBM Cloud App Management, Base and IBM Cloud App Management, Advanced (continued)

Title on Passport Advantage	eAssembly number	Product documentation
IBM InfoSphere Federation Server V10.1, Multilingual eAssembly for IBM SmartCloud Application Performance Management 7.7.0.1 Standard Edition for IBM Cloud App Management	CJ4JUEN	<a href="#">Click here</a>

## System requirements

Before installing the IBM Cloud App Management server, review the server hardware and software requirements. Before installing the ICAM Agents, or Kubernetes data collector review the prerequisites.

### IBM Cloud App Management hardware requirements

For details regarding CPU, RAM, VMs, and disk space requirements, see [“Planning hardware and sizing”](#) on page 77.

### IBM Cloud App Management server software requirements

For details on supported operating systems and platform, view the [IBM Cloud App Management detailed system requirements report](#).

For details on supported browsers, see the [Supported browsers](#) topic in the IBM Cloud Private Knowledge Center.

### ICAM Agents and data collectors

For details on the supported operating systems for each agent, review the following table. The table provides a link to the detailed system requirements report for each agent:

<i>Table 8. Supported operating systems for agents</i>	
Agents and data collectors	System requirements
Amazon EC2 agent	<a href="#">Amazon EC2 agent</a>
Amazon ELB agent	<a href="#">Amazon ELB agent</a>
Azure Compute agent	<a href="#">Azure Compute agent</a>
Cisco UCS agent	<a href="#">Cisco UCS agent</a>
Citrix VDI agent	<a href="#">Citrix VDI agent</a>
2019.4.0.2 CouchDB agent	<a href="#">CouchDB agent</a>
DataPower agent	<a href="#">DataPower agent</a>
DataStage agent	<a href="#">DataStage agent</a>
Db2 agent	<a href="#">Db2 agent</a>
Hadoop agent	<a href="#">Hadoop agent</a>
HMC agent	<a href="#">HMC agent</a>
Hyper-V Server agent	<a href="#">Hyper-V Server agent</a>
HTTP Server agent	<a href="#">HTTP Server agent</a>
JBoss agent	<a href="#">JBoss agent</a>
IBM Integration Bus agent	<a href="#">IBM Integration Bus agent</a>

<i>Table 8. Supported operating systems for agents (continued)</i>	
<b>Agents and data collectors</b>	<b>System requirements</b>
Linux OS agent	<a href="#">Linux OS agent</a>
Linux KVM agent	<a href="#">Linux KVM agent</a>
MariaDB agent	<a href="#">MariaDB agent</a>
Microsoft .NET agent	<a href="#">Microsoft .NET agent</a>
Microsoft Active Directory agent	<a href="#">Microsoft Active Directory agent</a>
Microsoft Cluster Server agent	<a href="#">Microsoft Cluster Server agent</a>
Microsoft Exchange Server agent	<a href="#">Microsoft Exchange Server agent</a>
Microsoft IIS agent	<a href="#">Microsoft IIS agent</a>
Microsoft Office 365 agent	<a href="#">Microsoft Office 365 agent</a>
Microsoft SharePoint Server agent	<a href="#">Microsoft SharePoint Server agent</a>
Microsoft SQL Server agent	<a href="#">Microsoft SQL Server agent</a>
MongoDB agent	<a href="#">MongoDB agent</a>
MySQL agent	<a href="#">MySQL agent</a>
NetApp Storage agent	<a href="#">NetApp Storage agent</a>
Oracle Database agent	<a href="#">Oracle Database agent</a>
PostgreSQL agent	<a href="#">PostgreSQL agent</a>
RabbitMQ agent	<a href="#">RabbitMQ agent</a>
SAP agent	<a href="#">SAP agent</a>
SAP HANA Database agent	<a href="#">SAP HANA Database agent</a>
SAP NetWeaver Java Stack agent	<a href="#">SAP NetWeaver Java Stack agent</a>
Skype for Business Server agent	<a href="#">Skype for Business Server agent</a>
<b>2019.4.0.2</b> Sterling Connect Direct agent	<a href="#">Sterling Connect Direct agent</a>
<b>2019.4.0.2</b> Sterling File Gateway agent	<a href="#">Sterling File Gateway agent</a>
Sybase agent	<a href="#">Sybase agent</a>
Tomcat agent	<a href="#">Software Tomcat agent</a>
UNIX OS agent	<a href="#">UNIX OS agent</a>
VMware VI agent	<a href="#">VMware VI agent</a>
WebLogic agent	<a href="#">WebLogic agent</a>
WebSphere Applications agent	<a href="#">WebSphere Applications agent</a>
WebSphere Infrastructure Manager agent	<a href="#">WebSphere Infrastructure Manager agent</a>
IBM MQ(formerly WebSphere MQ) agent	<a href="#">IBM MQ(formerly WebSphere MQ) agent</a>
Windows OS agent	<a href="#">Windows OS agent</a>
Kubernetes data collector	<a href="#">Kubernetes data collector</a>



<i>Table 8. Supported operating systems for agents (continued)</i>	
<b>Agents and data collectors</b>	<b>System requirements</b>
Runtime data collectors	<a href="#">Go data collector</a> <a href="#">Liberty data collector</a> <a href="#">J2SE data collector</a> <a href="#">Node.js data collector</a> <a href="#">Python data collector</a> <a href="#">Ruby data collector</a>
Synthetics PoP	<a href="#">Synthetics PoP</a>
Unified Agent plug-ins	<a href="#">UA plug-in for IBM API Connect</a> <a href="#">UA plug-in for IBM MQ</a> <a href="#">UA plug-in for Jaeger and Zipkin</a> <a href="#">UA plug-in for NGINX</a> <a href="#">UA plug-in for Redis</a> <a href="#">UA plug-in for IBM App Connect Enterprise</a> <a href="#">UA plug-in for DEM</a> <a href="#">UA plug-in for OpenShift</a>
DEM plug-ins	<a href="#">DEM plug-in for HTTP Server</a>

### Supported docker versions

IBM Cloud Private provides Docker packages that can be used for installation on boot and cluster nodes. This package is available for Red Hat Enterprise Linux and Ubuntu systems only. For more information, see the [Supported Docker Versions](#) topic in the IBM Cloud Private Knowledge Center.

### Supported file system and storage types

Storage class recommendations are local storage or vSphere. You are prompted to select one of these storage types during the IBM Cloud App Management installation. For more information see [Determine the storage type to use](#) in the *Planning your deployment* topic

## Planning hardware and sizing

You must allocate hardware in your Cloud App Management environment based on the number of monitored resources and the number of metrics that are uploaded per minute.

### Determine and specify your environment size in the `prepare-pv.sh` script

There are two sizing options for the container Kubernetes resource requests and limits. The following two options are available for sizing your Cloud App Management environment:

#### Demonstration/Proof of Concept

This size is suitable for a small demonstration, trial, or proof of concept. It is only suitable for a minimal workload. This size is designed to reduce the size of the microservices that are deployed to minimize the required hardware. Size0 is a minimum setting, which is not intended for horizontal scaling as the base per container overhead would be inefficient.

For a demonstration/proof of concept, specify Size0 when you run the `prepare-pv.sh` script prior to installing Cloud App Management, see [“3”](#) on page 78 below.

## Production

In a Production deployment, the container resources are larger than in a Demonstration/PoC environment. They have been vertically scaled for further scalability.


For stateless deployments that need additional capacity, use Horizontal Pod Autoscaler to increase the number of pods running. For more information, see [“Scaling stateless and stateful services”](#) on page 83 .

For statefulsets (Cassandra, Kafka, ZooKeeper, CouchDB, Datalayer, Elasticsearch), administrators will need to manually choose to increase the scale.

For a production implementation, specify Size1 when you run the prepare-pv.sh script prior to installing Cloud App Management, see [“3”](#) on page 78 below.

### How do you determine and specify what size to use?

To determine which size to use, complete the following steps:

1. Enter an estimate for the number of agents and data collectors in the projection spreadsheet, available here: [IBM Cloud App Management Load Projections Spreadsheet](#) . The spreadsheet returns an estimate of total metrics per minute. The 1 TB disk storage estimate for Cassandra in [Table 9 on page 78](#) is a conservative estimate. This spreadsheet gives you a more accurate estimate of the disk storage size required for Cassandra.
2. Review [Table 9 on page 78](#) to determine whether you should deploy a Size0 or a Size1 environment.
3. When you run the prepare-pv.sh script in preparation for the server installation, specify either Size0 or Size1. There are different parameters that you can choose depending on your platform and CPU requirements. Available sizes are:

```
--size0, --size1, --size0_amd64, --size0_ppc64le, --size_1_amd_64, --size1_ppc64le
```

For POWER9 processor, select **--size0\_ppc64le** or **--size1\_ppc64le**. This reduces the CPU allocation to reflect the generally lower CPU required by POWER9.

For POWER8 and earlier you do not need reduced CPU allocation and can select **size 0** or **size 1** parameter.

For AMD64, use **--size0\_amd64** or **--size\_1\_amd\_64**. Note these are effectively the same as **size 0** and **size 1**.

For detail of CPU requirement, see the Virtual Cores and Memory row in [Table 1](#) below.

The following table describes the hardware and minimum configuration for the different environment sizes.

Category	Resource	Demo/POC/Trial		Entry		Standard	Enterprise
Monitored environment size	Max metrics per minute	25,000	50,000	1,000,000	1,000,000	2,000,000	3,000,000
	Approx. resources	50	100	3,000	3,000	6,000	9,000
Environment options	Container size	Size0	Size0	Size1	Size1	Size1	Size1
	High Availability	No	Yes	No	Yes	Yes	Yes
Workers (OCP Compute) VMs	Minimum	1	3	1	3	6	9
	Recommended	2	3	3	6	9	13

Category	Resource	Demo/POC/Trial		Entry		Standard	Enterprise
Virtual cores and Memory	AMD (default) CPUs	10	20	35	70	90	105
	IBM Power CPUs	5	10	18	35	45	53
	Memory (GB)	32	64	55	130	180	230
Virtual cores and Memory with baselines enabled*	AMD (default) CPUs	12	24	40	75	95	110
	IBM Power CPUs	6	12	20	38	48	55
	Memory (GB)	32	72	65	140	195	250
Disk Storage	Cassandra	50 GB on 1 worker	100 GB on 3 workers	1 TB on 1 worker	1 TB on 3 workers	1 TB on 6 workers	1 TB on 9 workers
	Kafka	5 GB on 1 worker	10 GB on 3 workers	100 GB on 1 worker	100 GB on 3 workers	150 GB on 3 workers	200 GB on 4 workers
	CouchDB	5 GB on 1 worker	5 GB on 3 workers	25 GB on 1 worker	25 GB on 3 workers	50 GB on 3 workers	75 GB on 3 workers
	Datalayer	5 GB on 1 worker	5 GB on 3 workers	25 GB on 1 worker	25 GB on 3 workers	50 GB on 3 workers	75 GB on 3 workers
	Elasticsearch	5 GB on 1 worker	5 GB on 3 workers	25 GB on 1 worker	25 GB on 3 workers	50 GB on 3 workers	75 GB on 3 workers
	ZooKeeper	1 GB on 1 worker	1 GB on 3 workers	1 GB on 1 worker	1 GB on 3 workers	1 GB on 3 workers	1 GB on 3 workers

\*Baselines is not currently available

**Remember:**

The Virtual Cores and Memory row in Table 1 refers to total cores and memory across all workers. Allow 1 core for running each worker VM. For example, if Cloud App Management requires 24 cores and you're installing on 3 worker VMs, your total CPU requirement would be 24+3 cores.

Cassandra requires 4 cores, and 16GB available on the system. Given the IBM Cloud Private overhead, the minimum VM size possible to run on is 6 cores, and 20GB. We recommend running Cassandra on 8 cores, and 32GB or larger systems. If your disk IO is not sufficient, you can increase the RAM on your Cassandra and Cassandra containers to improve response time. For more information, see [“Cassandra Disk and Memory Usage - Adding Memory Resources”](#) on page 161.

In a production environment, as the number of agents, data collector, and metric traffic grows, you might need to horizontally scale your environment to handle the additional workload. Once the existing capacity of microservices can no longer handle the workload, more replicas can be deployed. For more information, see [“Scaling stateless and stateful services ”](#) on page 83.

Scaling out a non high-availability deployment increases the risk of an environment outage and data loss, as a single Cassandra disk failure will break the cluster's storage. This is why it is not recommended for larger deployments where a single Cassandra is insufficient.

Currently, 4,000 is the maximum supported resources of a single resource type. For example, you can have 3,000 Linux OS agents and 3000 UNIX OS agents, but not 6,000 Linux OS agents.

Cassandra requires a minimum of 3 nodes for high availability, as a "quorum" (3, tie-breaker) is required for critical data like topology or event records. More information about Cassandra requirements in a high availability environment are available in [“Planning for a high availability installation”](#) on page 148.

The minimum VMs recommendation in the Workers vms row of Table 1 is based on spreading the statefulset (database) data to different systems to minimize the risk of an outage. Generally, the minimum number of required VMs will match the number of Cassandra replicas needed. Depending on the size of your individual VMs, you might need more systems to reach the total CPU and memory required to deploy the Cloud App Management pods. The recommend distribution includes additional VMs for the other statefulsets, however these can be placed on the same node as Cassandra if it has the CPU, memory, and Disk capacity. Review information for deploying stateful services in a high availability environment in the [“Planning for a high availability installation”](#) on page 148 topic.

For data collectors, keep in mind that the number of pods that you are monitoring with a single data collector can be considerable and might warrant a Size1 environment.

Network based storage like NFS or Gluster is not recommended, as the disk IO of Cassandra can easily saturate any network attached or mounted device.

## Planning ports

Before you begin your deployment of IBM Cloud Private, and Cloud App Management, you must consider what ports to open.

### IBM Cloud Private

Before you install IBM Cloud Private, review the information in the following IBM Cloud Private knowledge center topic: [Default ports](#).

### Cloud App Management ports

Before you install Cloud App Management, you must open the default ports: 443 and 8080.

### ICAM Agents

Before you install the ICAM Agents, review the following information to determine which ports to open:

Agents	Default ports	Configurable
Amazon EC2 agent	<ul style="list-style-type: none"> <li>TCP port 80 (for HTTP)</li> <li>TCP port 443 (for HTTPS)</li> </ul>	N/A
Amazon ELB agent	<ul style="list-style-type: none"> <li>TCP port 80 (for HTTP)</li> <li>TCP port 443 (for HTTPS)</li> </ul>	N/A
Azure Compute agent	<ul style="list-style-type: none"> <li>TCP port 80 (for HTTP)</li> <li>TCP port 443 (for HTTPS)</li> </ul>	N/A
Cassandra agent	7199 (for JMX server, local and remote)	Yes
Cisco UCS agent	N/A	N/A
Citrix VDI agent	5986 (for Power Shell)	Yes
<b>2019.4.0.2</b> CouchDB agent	5984	Yes
DataPower agent	5550 (for connecting to remote DataPower appliances)	Yes

Table 10. Default ports used by agents (continued)

Agents	Default ports	Configurable
Db2 agent	N/A	N/A
DataStage agent	<ul style="list-style-type: none"> <li>• 9443 (WAS HTTPS port on Windows)</li> <li>• 9446 (WAS HTTPS port on Linux)</li> <li>• 50000 (Database JDBC port for DB2)</li> <li>• 1433 (Microsoft SQL)</li> <li>• 1521 (Oracle)</li> </ul>	Yes
HMC agent	22	
Hadoop agent	<ul style="list-style-type: none"> <li>• Local monitoring: CP_PORT environment variable value</li> <li>• Remote monitoring:                             <ul style="list-style-type: none"> <li>– 50070 (Standby Namenode)</li> <li>– 50090 (Secondary Namenode)</li> <li>– 8088 (ResourceManager)</li> <li>– 19888 (JobHistory Server)</li> <li>– 8080 (Ambari)</li> </ul> </li> </ul>	Yes
IBM Integration Bus agent	N/A	N/A
Linux KVM agent	<ul style="list-style-type: none"> <li>• 8080 (for HTTP)</li> <li>• 8443 (for HTTPS)</li> </ul>	Yes
JBoss agent	N/A	Yes
MariaDB agent	3306	Yes
Microsoft .NET agent	N/A	N/A
Microsoft Active Directory agent	The port number depends on the listener setting for monitoring usage.	No
Microsoft Cluster Server agent	N/A	N/A
Microsoft Exchange Server agent	N/A	N/A
Microsoft IIS agent	N/A	N/A
Microsoft Office 365 agent	N/A	N/A
Microsoft SharePoint Server agent	N/A	N/A
Microsoft SQL Server agent	N/A	N/A
MongoDB agent	<ul style="list-style-type: none"> <li>• 27017 (for single instance)</li> <li>• 27019 (for cluster)</li> </ul>	Yes
MySQL agent	3306 (for JDBC connection)	Yes

Table 10. Default ports used by agents (continued)

Agents	Default ports	Configurable
NetApp Storage agent	For remote monitoring: <ul style="list-style-type: none"> <li>• 8088</li> <li>• 8488</li> <li>• 443</li> <li>• 8443</li> </ul>	No
Oracle Database agent	1521 (for SQL connection)	Yes
PostgreSQL agent	5432 (for JDBC connection)	Yes
RabbitMQ agent	15672	No
UNIX OS agent	N/A	N/A
SAP agent	33nn (where nn is the SAP instance number)	Yes
SAP HANA Database agent	<ul style="list-style-type: none"> <li>• Default: 30013</li> <li>• Range: 30013-39913</li> <li>•</li> </ul>	Yes
SAP NetWeaver Java Stack agent	<ul style="list-style-type: none"> <li>• Default: 50004</li> <li>• Range: 50004-59904</li> </ul>	Yes
Skype for Business Server agent	5061	N/A
<b>2019.4.0.2</b> Sterling Connect Direct agent	1363	Yes
<b>2019.4.0.2</b> Sterling File Gateway agent	50000 The IBM B2B Integrator REST API port number and Database port number are both required and are configurable.	Yes
Sybase agent	5000	No
Tomcat agent	<ul style="list-style-type: none"> <li>• 8686 (JMX port)</li> <li>• 8080 (default)</li> </ul>	Yes
VMware VI agent	<ul style="list-style-type: none"> <li>• 443 (remote monitoring)</li> <li>• 80 (local monitoring)</li> </ul>	Yes
WebLogic agent	7003	Yes
WebSphere Applications agent	63355 (for resource monitoring)	Yes
WebSphere Infrastructure Manager agent	N/A	N/A
IBM MQ(formerly WebSphere MQ) agent	N/A	N/A
Windows OS agent	N/A	N/A

## Scaling stateless and stateful services

As metric traffic grows and your agent numbers increase, you must consider horizontally scaling your environment. Scaling your environment means adding additional services. In the context of scaling, there are two types of services: stateless services and stateful services.

### Stateless services

The stateless services in Cloud App Management are automatically scaled using Horizontal Pod Autoscaler (HPA). HPA scales up and down the number of replicas based on the CPU usage of the service. By default, the HPA upscale-delay is 3 minutes, and HPA downscale delay is five minutes. This means Kubernetes will wait for usage to stabilize for three minutes before scaling up and five before scaling down. Each HPA has a threshold value, which is compared against the CPU request. For example, if the HPA threshold is 80% and the CPU request is 1000m (1 CPU core), the HPA will trigger a scale up after running at or above 800m for 3 minutes. Stateless services are listed here in the reference section.

While the HPA provides scalability for stateless services to provide high availability for services, you need to take some extra steps. The HPAs have a minimum and maximum number of replicas they can scale to. By adjusting the `--minReplicasHPAs` parameter to a number greater than one, you can provide a level of high availability by ensuring that multiple instances are deployed by Kubernetes. You set these parameters when you run the `pre-install.sh` script. For more information, see [“Planning for a high availability installation” on page 148](#).

To manually edit HPAs, use the following command:

```
kubectl edit hpa releasename-service
```

For example

```
kubectl edit hpa ibmcloudappmgmt-metric
```

Alternatively, to edit HPAs in the IBM Cloud Private UI, from the menu, select Configuration>Scaling Policies, select the policy and edit the service. For more information, see the [Horizontal pod auto scaling by using custom metrics](#) in the IBM® Cloud Private Knowledge Center.

### Stateful services

You must manually scale stateful services:

For Cassandra, see [“Scaling Cassandra ” on page 84](#).

For Kafka, see [“Scaling up Kafka brokers” on page 86](#).

For CouchDB, see [“Scaling up CouchDB” on page 85](#).

For Datalayer, see [“Scaling up Datalayer” on page 87](#).

For ZooKeeper, see [“Scaling up ZooKeeper” on page 87](#).

### List of stateless services that can scale with HPA

- agentbootstrap
- agentmgmt
- alarmeventsrc
- amui
- applicationmgmt
- config
- event-observer
- ibm-cem-rba-as
- ibm-cem-brokers
- ibm-cem-cem-users
- ibm-cem-channelservices

- ibm-cem-event-analytics-ui
- ibm-cem-eventpreprocessor
- ibm-cem-incidentprocessor
- ibm-cem-integration-controller
- ibm-cem-normalizer
- ibm-cem-notificationprocessor
- ibm-cem-rba-rbs
- ibm-cem-scheduling-ui
- linking
- metric
- metricenrichment
- metricprovider
- opentt-collector
- opentt-query
- synthetic
- temacomm
- temaconfig
- temasda
- threshold

## Scaling Cassandra

The Cassandra StatefulSet provides the incident store.

### About this task

Cassandra requires 1 node per 1 million metrics per minute. This metric equates to approximately Linux OS agent x 4000, or WebSphere Applications agent x 500. To estimate your metrics per minute, enter an estimate of the number of agents into the projections spreadsheet. For more information, see [“Planning hardware and sizing”](#) on page 77.

If the metric data replication factor is set to 3 (default for high availability environments), you will need 3 Cassandra nodes per 1 million metrics per minute. For example, if you expect 3 million metrics per minute and run with metric replication of 3, you will need at least 9 Cassandra nodes. In highly available environments, you are allowed to expand Cassandra one node at a time. Cassandra distributes the workload evenly to all members of the cluster. If you have 4 nodes with replication factor of 3, each node will hold roughly 75% of the total data.

Cassandra requires 4 cores, and 16GB available on the system. Given the IBM Cloud Private overhead, the minimum VM size possible to run on is 6 cores, and 20GB. We recommend running Cassandra on 8 cores, and 32GB or larger systems.

Check all pods are ready before adding another pod. Don't increase the number of pods by more than one at a time; wait until the new node is ready before adding more. Adding nodes can cause a large amount of network traffic, so schedule it at a quieter time.

**Tip:** To optimize performance in the resource dashboards, see [“Optimizing disk performance for Cassandra”](#) on page 156.

### Procedure

1. If you're using local storage, you must add an extra persistent volume for Cassandra. The number of StatefulSets starts at 0.

Run the following two commands:



```
./ibm_cloud_pak/pak_extensions/lib/cloud-pv.sh \
--release my_release name \
--name my_release name-cassandrapv_number \ # Add the PV number immediately after -cassandra
--node my_node \
--class my_release name-local-storage-cassandra \
--dir my_directory \
--size 2000Gi
```

For example,

```
./ibm_cloud_pak/pak_extensions/lib/cloud-pv.sh \
-- release ibmcloudappmgmt \
-- name ibmcloudappmgmt-cassandra1 \
-- node perfvm4123 \
-- class ibmcloudappmgmt-local-storage-cassandra \
-- dir /k8s/data/cassandra \
--size 2000Gi \
```

Run the following command:

```
kubectl create -f ./ibm-cloud-appmgmt-prod/yaml/PersistentVolume_my_release name
-cassandra1_my_release_name.yaml
```

For example,

```
kubectl create -f ./ibm-cloud-appmgmt-prod/yaml/PersistentVolume_ibmcloudappmgmt
-cassandra1_ibmcloudappmgmt.yaml
```

2. Increase the Cassandra StatefulSet scale count using either of the following methods:

- In the IBM Cloud Private UI, select **Workloads>StatefulSets**, select the appropriate Cassandra, select **Action>Scale** and enter the increased count.
- Using the CLI, scale the StatefulSet, for example:

```
kubectl scale --replicas=2 statefulset/ibmcloudappmgmt-cassandra
```

3. Spread data to the new Cassandra replica. Space on the existing nodes won't be reclaimed immediately. Use the following command on one pod at a time to reclaim the space:

```
kubectl exec -i my_release_name-cassandra-0 --
/opt/ibm/cassandra/bin/nodetool cleanup
```

## Scaling up CouchDB

You can increase the number of CouchDB pods during an upgrade or installation by increasing the replicas value of the CouchDB Stateful Set -couch db.

### About this task

After installation, you can increase the number of CouchDB pods by increasing the replicas value of the CouchDB StatefulSet -couchdb by either using the IBM Cloud Private UI console or completing the following steps:

### Procedure

1. If you're using local storage, you must add an extra persistent volume for CouchDB. The number of StatefulSets starts at 0. Run the following command:

```
./ibm_cloud_pak/pak_extensions/lib/cloud-pv.sh \
--release my_release_name \
--name my_release_name-couchdbpv_number \ ## Add the PV number immediately
after -couchdb
--node my_node \
--class my_release_name-local-storage-couchdb \
--dir my_directory \
--size 1Gi
```

For example,

```
./ibm_cloud_pak/pak_extensions/lib/cloud-pv.sh \
-- release ibmcloudappmgmt \
-- name ibmcloudappmgmt-couchdb\
-- node worker04 \
-- class ibmcloudappmgmt-local-storage-couchdb\
-- dir /k8s/data/couchdb\
--size 1Gi \
```

2. Use the Kubernetes scale command to scale replicas, for example:

```
kubectl scale sts releasename-couchdb --replicas=3
```

## Scaling up Kafka brokers

In a horizontally scaled environment, you might need to manually add additional Kafka brokers.

### Procedure

1. If you're using local storage, you must add an extra persistent volume for Kafka. The number of StatefulSets starts at 0. Run the following command:

```
./ibm_cloud_pak/pak_extensions/lib/cloud-pv.sh \
--release my_release_name \
--name my_release_name-kafkapv_number \# Add the PV number immediately after -kafka
--node my_node \
--class my_release_name-local-storage-kafka \
--dir my_directory \
--size 1Gi
```

For example,

```
./ibm_cloud_pak/pak_extensions/lib/cloud-pv.sh \
-- release ibmcloudappmgmt \
-- name ibmcloudappmgmt-kafka \
-- node worker04 \
-- class ibmcloudappmgmt-local-storage-kafka \
-- dir /k8s/data/kafka \
--size 1Gi
```

2. To scale up the number of brokers, for example, to scale up from 3 to 6 , complete the following steps:
  - a) In one command window, run the following command to watch the changes to the StatefulSet:

```
kubectl get pods -w -l app=releasename-kafka
```

- b) In another command window, run the following command to increase the number of Kafka brokers:

```
kubectl scale sts releasename-kafka --replicas=6
```

where *releasename* is the name of the Kafka broker

3. Once the Pods are available, decide which topics need to be reassigned, these are the topics managed by Cloud App Management server:
4. Access one of the Kafka Pods by running the following command:

```
kubectl exec -it releasename-kafka-0 bash
```

5. Create a file with the topics whose partitions you wish to reassign, for example:

```
cat >/tmp/topics-to-move.json <<EOF
{
  "topics": [
    {"topic": "incidents"},
    {"topic": "cem-notifications"}
  ],
  "version": 1
}
EOF
```

6. Run the following command to get the list of brokers:

```
/opt/kafka/bin/zookeeper-shell.sh $ZOOKEEPER_URL <<< "ls /brokers/ids"
```

7. On the Kafka Pod, run the following command with the list of brokers obtained from the step “6” on page 86, for example :

```
/opt/kafka/bin/kafka-reassign-partitions.sh --topics-to-move-json-file  
/tmp/topics-to-move.json --broker-list "0,1,2,3,4,5" --zookeeper $ZOOKEEPER_URL  
--generate | grep version | grep partitions | tail -1 >/tmp/new-replicas.json
```

where *new-replicas.json* is the name of the new json file created

8. Review the *new-replicas.json* file and make modifications as required.
9. Run execute on the *new-replicas.json* file, for example:

```
/opt/kafka/bin/kafka-reassign-partitions.sh --reassignment-json-file  
/tmp/new-replicas.json --zookeeper $ZOOKEEPER_URL --execute
```

## Scaling up ZooKeeper

You can increase the number of ZooKeeper pods during an upgrade or installation by increasing the replicas value of the ZooKeeper Stateful Set `-zookeeper`

### About this task

After installation, you can increase the number of ZooKeeper pods by increasing the replicas value of the `zookeeperStatefulSet -zookeeper` by either using the IBM Cloud Private UI console or completing the following steps:

### Procedure

1. If you're using local storage, you must add an extra persistent volume for Zookeeper. The number of StatefulSets starts at 0. Run the following command:

```
./ibm_cloud_pak/pak_extensions/lib/cloud-pv.sh \  
  --release my_release_name \  
  --name my_release_name-zookeeper1pv_number \  
after -zookeeper \  
  --node my_node \  
  --class my_release_name-local-storage-zookeeper \  
  --dir my_directory \  
  --size 1Gi
```

For example,

```
./ibm_cloud_pak/pak_extensions/lib/cloud-pv.sh \  
  -- release ibmcloudappmgmt \  
  -- name ibmcloudappmgmt-zookeeper1\  
  -- node worker04 \  
  -- class ibmcloudappmgmt-local-storage-zookeeper\  
  -- dir /k8s/data/zookeeper\  
  --size 1Gi \  
  --
```

2. Use the Kubernetes scale command to scale replicas, for example:

```
kubectl scale sts releasename-zookeeper --replicas=3
```

## Scaling up Datalayer

You can increase the number of Datalayer pods during an upgrade or installation by increasing the replicas value of the Datalayer stateful set `-couchdb`

### About this task

## Procedure

1. If you're using local storage, you must add an extra persistent volume for Datalayer. The number of StatefulSets starts at 0. Run the following command:

```
./ibm_cloud_pak/pak_extensions/lib/cloud-pv.sh \  
    --release my_release_name \  
    --name my_release_name-datalayerpv_number \  
after -datalayer \ # Add the PV number immediately  
    --node my_node \  
    --class my_release_name-local-storage-datalayer \  
    --dir my_directory \  
    --size 1Gi
```

For example,

```
./ibm_cloud_pak/pak_extensions/lib/cloud-pv.sh \  
    -- release ibmcloudappgmt \  
    -- name ibmcloudappgmt-datalayer\  
    -- node worker04 \  
    -- class ibmcloudappgmt-local-storage-datalayer\  
    -- dir /k8s/data/datalayer\  
    --size 1Gi \
```

2. Use the Kubernetes scale command to scale replicas, for example:

```
kubectl scale sts releasename-datalayer --replicas=3
```

---

## Chapter 8. Deploying IBM Cloud Private

You must install and deploy the IBM Cloud Private platform before you can install and deploy the IBM Cloud App Management product. IBM Cloud Private is an application platform for developing and managing on-premises containerized applications. You can download the installation package for IBM Cloud Private Enterprise, V3.2.1 from [IBM Passport Advantage](#).

### Preparing to install an IBM Cloud Private Enterprise environment

Before you install IBM Cloud Private Enterprise version 3.2.1, you must prepare your system. Preparation tasks include; checking system requirements, choosing storage and port options, configuring your cluster, and installing Docker.

#### About this task

IBM Cloud Private Enterprise supports the Linux 64-bit platform. The Cloud App Management and IBM Cloud Private Enterprise products support the following operating systems:

- Red Hat Enterprise Linux (RHEL) versions 7.3 x86-64, 7.4 x86-64, and 7.5 x86-64
- Ubuntu 16.04 LTS

#### Procedure

1. Review the Cloud App Management [“System requirements” on page 75](#) and the IBM Cloud Private Enterprise [System requirements](#). Ensure that your system meets these requirements before you install the IBM Cloud Private Enterprise platform.
2. Prepare your cluster:
  - a. Determine your cluster architecture, and obtain the IP address for all nodes in your cluster. For more information about node types, see the [Architecture](#) topic in the IBM Cloud Private IBM Knowledge Center.

**Note:** During IBM Cloud Private installation, you add the IP address for your master, worker, and proxy nodes to this file. You can also specify a management node. For details, see the [Setting the node roles in the hosts file](#) topic in the IBM Cloud Private IBM Knowledge Center. After you install IBM Cloud Private Enterprise, you can add or remove management, worker, and proxy nodes from your cluster.

**Note:** IBM Cloud Private requires Docker. During the Docker installation, be sure to review the [Configuring your Docker engine](#) topic and set up Docker log rotation. This reduces disk issues that are caused by retaining too much log information.

- b. Prepare each node for installation. For more information, see the [Configuring your cluster](#) topic in the IBM Cloud Private IBM Knowledge Center.
3. Optional: Add extra hard disk drives on all your nodes for local storage. For more information about file storage, see the [Supported file systems and storage](#) topic in the IBM Cloud Private IBM Knowledge Center. If you have VMWare and vSphere administrator access, you can use the extra drive for your storage class. For more details, see the [Creating a storage class for vSphere volume](#) topic in the IBM Cloud Private IBM Knowledge Center.

**Note:** A recommendation is to add one or more extra drives to store persistent data that is needed for Cloud App Management, Apache Cassandra, Apache Kafka, Apache ZooKeeper, CouchDB, and the IBM Cloud Event Management data layer. If you are running a small environment, you should have enough space on the / drive. However, for a larger scale environment, you might need TBs of storage space, which would mostly be used for storing metric data in Cassandra.

4. Optional: If you want to use a different default Docker storage directory, you must change it before you install IBM Cloud Private by using a bind mount. For more information, see the [Specify a default Docker storage directory by using bind mount](#) topic in the IBM Cloud Private IBM Knowledge Center.

5. Optional: If you want to use different default storage directories for the core IBM Cloud Private services, you must change them before you install IBM Cloud Private by using a bind mount. For more information, see the [Specify other default storage directories by using bind mount](#) topic in the IBM Cloud Private IBM Knowledge Center.

## Installing an IBM Cloud Private Enterprise environment

Learn how to successfully download and install IBM Cloud Private Enterprise V 3.2.1.

### Before you begin

You must complete all the prerequisite steps in the [“Preparing to install an IBM Cloud Private Enterprise environment”](#) on page 89 topic.

### About this task

For more information about installing IBM Cloud Private Enterprise, see the [Installing IBM Cloud Private Cloud Native and Enterprise editions](#) topic in the IBM Cloud Private IBM Knowledge Center.

### Procedure

1. Manually install Docker on your boot node. For more information about the boot node, see the [Boot node](#) topic in the IBM Cloud Private IBM Knowledge Center. For more information about manually installing Docker on your boot node, see the [Setting up Docker for IBM Cloud Private](#) topic in the IBM Cloud Private IBM Knowledge Center.

If you already installed Docker on your boot node, ensure that the version is supported. For a list of Docker versions that are supported by IBM Cloud Private, see the [Supported Docker versions](#) topic in the IBM Cloud Private IBM Knowledge Center.

**Note:** During the Docker installation, be sure to review the [Configuring your Docker engine](#) topic and set up Docker log rotation. This reduces disk issues that are caused by retaining too much log information.

2. Set up the installation environment. Download the installation files for IBM Cloud Private from the [IBM Passport Advantage](#) website. You must download the correct file or files for the type of nodes in your cluster. To set up the installation environment, complete all steps in the *Set up the installation environment* section of the [Installing an IBM Cloud Private Enterprise environment](#) topic in the IBM Cloud Private IBM Knowledge Center.
3. Optional: To customize your cluster, complete all steps in the *Customize your cluster* section of the [Installing an IBM Cloud Private Enterprise environment](#) topic in the IBM Cloud Private IBM Knowledge Center.
4. Install Docker on your cluster nodes. IBM Cloud Private can install Docker on cluster nodes as part of the installation process. However, to speed up the IBM Cloud Private deployment, manually installing Docker on each node and populating each Docker image repository is recommended. The IBM Cloud Private Docker binary package is available to use for Docker installation. For more information, see the [Supported Docker Versions](#) topic and the *Manually installing Docker by using the provided IBM Cloud Private Docker package* section on the [Setting up Docker for IBM Cloud Private](#) topic in the IBM Cloud Private IBM Knowledge Center.
  - a. For Red Hat Enterprise Linux (RHEL) systems, the storage driver for the supplied IBM Cloud Private Docker package is set to `loop-lvm` by default. For production deployments, you must change to a different storage option. Configure either `direct-lvm` or `overlay2` in your production environment.

**Note:** The `overlay2` storage driver can be configured on Red Hat Enterprise Linux (RHEL) V7.5 or higher systems. For more information, see *Docker storage drivers* in the [Docker](#) documentation. The system must use an xfs formatted drive. Confirm that your device can run the `overlay2`

storage driver, by running the **xfs\_info** command. Check for the `fstype=1` value in the output, as shown in the following example:

```
xfs_info /dev/sda1 | grep fstype
naming      =version 2          bsize=4096   ascii-ci=0  fstype=1
```

- Set up `direct-lvm`, as shown in the following example:

```
mkdir -p /etc/docker
cat >/etc/docker/daemon.json <<EOF
{
  "storage-driver": "devicemapper",
  "storage-opts": [
    "dm.directlvm_device=/dev/my_device",
    "dm.thinp_percent=95",
    "dm.thinp_metapercent=1",
    "dm.thinp_autoextend_threshold=80",
    "dm.thinp_autoextend_percent=20",
    "dm.directlvm_device_force=false"
  ]
}
EOF
```

Where `my_device` is the name of a required extra empty device.

- Set up `overlay2`, as shown in the following example:

```
mkdir -p /etc/docker
cat >/etc/docker/daemon.json <<EOF
{
  "storage-driver": "overlay2",
  "storage-opts": [
    "overlay2.override_kernel_check=true"
  ]
}
EOF
```

- b. After you set up the `etc/docker/daemon.json` file for either the `direct-lvm` or `overlay2` storage driver, install Docker and reconfigure it to use a storage driver other than the default, as shown in the following example:

```
./icp-docker-17.12.1_x86_64.bin --install
sed -i -e 's|ExecStart=/usr/bin/dockerd .*$|ExecStart=/usr/bin/dockerd --log-opt max-size=50m --log-opt max-file=10|' /usr/lib/systemd/system/docker.service
systemctl daemon-reload
systemctl restart docker
systemctl enable docker
```

- c. After you install Docker on your cluster nodes, run the following commands to verify your Docker installation:

```
docker info
systemctl status docker
```

5. Deploy the environment. For more information, see the *Deploy the environment* section of the [Installing an IBM Cloud Private Enterprise environment](#) topic in the IBM Cloud Private IBM Knowledge Center.
6. Access your IBM Cloud Private cluster. Log in to the IBM Cloud Private management console with a web browser. Go to your cluster URL, `https://my_master_ip:8443`, where `my_master_ip` is the IP address of the master node for your IBM Cloud Private cluster. Enter your login credentials. The default username is `admin`, and the default password is `admin`. This information is displayed in the installation logs, as shown in the following code:

```
UI URL is https://my_master_ip:8443, default username/password is admin/admin
```

7. Access the Docker private image registry. Configure authentication from your computer to the Docker private image registry host and login to the Docker private image registry. For more information, see

the [Configuring authentication for the DockerCLI](#) topic in the IBM Cloud Private IBM Knowledge Center .

### What to do next

When the IBM Cloud Private installation completes, complete the following tasks.

- IBM Cloud Private is affected by a privilege escalation vulnerability in the Kubernetes API server. To fix this issue and upgrade Kubernetes, download and apply the patch appropriate to your version from IBM Fix Central. For more information, see the [IBM Security Bulletin](#).
- If you previously disabled firewalls, restart your firewall.
- Ensure that all the default ports are open. For more information about ports, see [“Planning ports” on page 80](#).
- Back up the boot node. Copy your `/my_installation_directory/cluster` directory to a secure location. If you use SSH keys to secure your cluster, ensure that the SSH keys in the backup directory remain in sync.
- If you want to create more IBM Cloud Private users, complete the following steps:
  - If you have an LDAP directory, you can connect it with your IBM Cloud Private cluster. You can import users from your LDAP directory to add to your cluster. For more information, see the [Configuring LDAP connection](#) topic in the IBM Cloud Private IBM Knowledge Center.
  - If you want to create teams, add users to a team, or add groups to a team, see the [Teams](#) topic in the IBM Cloud Private IBM Knowledge Center.
- Deploy the Cloud App Management server and agents. For more information, see the following topics: [“Offline: Installing IBM Cloud App Management stand-alone on IBM Cloud Private” on page 148](#), [Chapter 13, “Deploying ICAM Agents,” on page 193](#), and [“Starting the Cloud App Management UI” on page 176](#).

## IBM Cloud Private postinstallation manual tasks

After you install IBM Cloud Private Enterprise, several tools must be installed before you can install the Cloud App Management server: IBM Cloud Private command line interface (CLI), `kubectl`, and Helm CLI.

You can install the IBM Cloud Private V3.2.1 command line interface (CLI) from the IBM Cloud Private management console. Older versions of the CLI don't work with IBM Cloud Private V3.2.1. You can also install the Kubernetes CLI and Helm CLI from the IBM Cloud Private console.

To install a CLI, click **Menu > Command Line Tools** and select the CLI you want to install. For more information, see the following topics in the IBM Cloud Private Knowledge Center:

- [Installing the IBM Cloud Private CLI](#)
- [Accessing your IBM Cloud Private cluster by using the kubectl CLI](#)
- [Setting up the Helm CLI](#)



# Chapter 9. Installing IBM Cloud App Management - the options

You have multiple different options to choose from when you are installing IBM Cloud App Management: install Cloud App Management with IBM Cloud Pak for Multicloud Management, install Cloud App Management standalone on Red Hat OpenShift, or on install Cloud App Management standalone on IBM Cloud Private.

Two types of installation methods are available:

### Online

Use Entitled Registry to pull the Cloud App Management image from the IBM Cloud entitled registry. The Helm chart is loaded in the IBM entitled charts in the catalog where it can be installed.

### Offline

The PPA (Passport Advantage Archive) file must be downloaded from IBM Passport Advantage. The images are extracted from it, and installed.

The following tables lists the installation options that are available for each method.

Online	Offline	Supported platforms and software required
<ul style="list-style-type: none"> <li>• <a href="#">“Online proof-of-concept installation of IBM Cloud App Management with IBM Cloud Pak for Multicloud Management” on page 96</a></li> <li>• <a href="#">“Online full monitoring installation of IBM Cloud App Management with IBM Cloud Pak for Multicloud Management” on page 99</a></li> </ul>	<ul style="list-style-type: none"> <li>• <a href="#">“Offline eventing only installation of IBM Cloud App Management with IBM Cloud Pak for Multicloud Management” on page 109</a></li> <li>• <a href="#">“Upgrading IBM Cloud App Management with IBM Cloud Pak for Multicloud Management from Eventing only to full monitoring mode” on page 181</a></li> <li>• <a href="#">“Offline full monitoring installation of IBM Cloud App Management with IBM Cloud Pak for Multicloud Management” on page 120</a></li> </ul>	<ul style="list-style-type: none"> <li>• Linux® x86_64 with Red Hat OpenShift V3.11</li> <li>• Linux® x86_64 with Red Hat OpenShift V4.2</li> <li>• Linux® on Power® with Red Hat OpenShift V3.11</li> </ul>
	<a href="#">“Offline: Installing IBM Cloud App Management stand-alone on Red Hat OpenShift” on page 141</a>	<ul style="list-style-type: none"> <li>• Linux® x86_64 with Red Hat OpenShift V3.11 and IBM Cloud Private V3.2.1 and its latest fix pack.</li> </ul>
	<a href="#">“Offline: Installing IBM Cloud App Management stand-alone on IBM Cloud Private” on page 148</a>	<ul style="list-style-type: none"> <li>• Linux® x86_64 with IBM Cloud Private V3.2.1 and its latest fix pack.</li> </ul>

### Related concepts

[“Troubleshooting installation and upgrade” on page 1395](#)

Troubleshoot IBM Cloud App Management installation and upgrade issues.

## Installing IBM Cloud App Management with IBM Cloud Pak for Multicloud Management

You can install IBM Cloud App Management with IBM Cloud Pak for Multicloud Management. With this combination, you can monitor cloud and on-premises application environments with IBM Cloud App Management and use IBM Cloud Pak for Multicloud Management to ensure that your clusters are secure, operating efficiently, and delivering expected service levels. IBM Cloud Pak for Multicloud Management provides user visibility and policy-based compliance across clouds and clusters. Two types of installation methods are available: online, which uses Entitled Registry to pull the Cloud App Management Passport Advantage Archive (PPA) installation image from the IBM Cloud entitled registry. The Helm chart is loaded in the IBM entitled charts in the catalog where it can be installed. The other method is offline, where you download the PPA (Passport Advantage Archive) installation image file from [IBM Passport Advantage®](#), extract the Helm chart from it, load it into the catalog, and continue with your installation.

### About this task

1. [“Onboarding LDAP users” on page 94](#)
2. Choose the IBM Cloud App Management with IBM Cloud Pak for Multicloud Management installation option that you want by reviewing the list in the [Chapter 9, “Installing IBM Cloud App Management - the options,” on page 93](#) topic.
3. Installing the ICAM klusterlet on the managed with or without the Helm:
  - a. [“Installing the ICAM klusterlet on the managed cluster with Helm” on page 133](#)
  - b. [“Installing the ICAM klusterlet on the managed cluster without helm” on page 135](#)
4. [“Deploying agents and data collectors for IBM Cloud Pak for Multicloud Management” on page 140.](#)
5. [“Uninstalling IBM Cloud App Management” on page 131.](#)

## Onboarding LDAP users

Before you install the Cloud App Management server into IBM Cloud Pak for Multicloud Management, import LDAP users, create an account, onboard users, set up a team, and add your managed cluster to the team.

### Before you begin

You must set up an LDAP connection in IBM Cloud Pak for Multicloud Management. From the navigation menu, select **Administer>Identity & Access**. Select **Create Connection**. The "LDAP Connection" page is displayed. For more information, see [Configuring LDAP connection](#).

**Required user type or access level:** Cluster administrator

### Procedure

Import users

1. Import users from the LDAP connection into IBM Cloud Pak for Multicloud Management:

```
cloudctl iam user-import -c LDAPID -u USERID
```

This step imports the LDAP user in to IBM Cloud Pak for Multicloud Management. Repeat this command for each LDAP user.

You can use this command to find the LDAP IDs:

```
cloudctl iam ldaps
```

Create an account

2. You must be logged in to IBM Cloud Pak for Multicloud Management as a cluster administrator to create an account. Create a user account in IBM Cloud Pak for Multicloud Management:

```
cloudctl login
```

Create an account:

```
cloudctl iam account-create NAME [-d, --description DESCRIPTION]
OPTIONS:
  -d, --description Description of the account
```

An account ID is returned. You use this *ACCOUNT\_ID* in a following step.

Onboard LDAP users

3. You must be logged in to IBM Cloud Pak for Multicloud Management as a cluster administrator to onboard LDAP users.

LDAP users can be onboarded with either the *PRIMARY\_OWNER* role or *MEMBER* role. Onboard at least one LDAP user with the *PRIMARY\_OWNER* role and the other remaining users with the *MEMBER* role. The user that is onboarded with the *PRIMARY\_OWNER* role takes on account administrator privileges and can log in to IBM Cloud Pak for Multicloud Management. The users who are added with the *MEMBER* role cannot log in to IBM Cloud Pak for Multicloud Management until they are added to a team in step 7.

Onboard the LDAP users imported in step “1” on page 94 to the account created in step 2:

```
cloudctl iam user-onboard ACCOUNT_ID -r accountRole -u user1ID,user2ID,...
OPTIONS:
  --role value, -r value Account role for user (PRIMARY_OWNER or MEMBER)
  -u value, --users value User or list of users to onboard onto account
```

**Note:** An LDAP user can be onboarded to only ONE account. If a user is mistakenly onboarded to multiple accounts, delete the user and onboard the user again.

Create a namespace.

4. You must be logged in to IBM Cloud Pak for Multicloud Management as a cluster administrator. From the IBM Cloud Pak for Multicloud Management console, select **Create resource**. Add a yaml or JSON file with the namespace details. You can also use the `kubectl create namespace` command to create a namespace by using the Kubernetes CLI. See [Installing the Kubernetes CLI \(kubectl\)](#) for instructions about how to install the Kubernetes CLI. For more information about creating a namespace, see [Creating a namespace](#).

Create a team.

5. Create a team for the *ACCOUNT\_ID* you created in step 2:

```
cloudctl iam team-create NAME
```

Add users to the team.

6. Add the users that were onboarded with the *MEMBER* role in step “3” on page 95 in to the team:

```
cloudctl iam team-add-users TEAM_ID ROLE -u user2ID
For example:
cloudctl iam team-add-users myteam Administrator -u user2ID
```

Add your managed cluster namespace in to the team you created.

7. You must be logged in to IBM Cloud Pak for Multicloud Management as a cluster administrator. From the IBM Cloud Pak for Multicloud Management navigation menu, select **Administer>Identity and Access>Teams**. Select the team that you created in step 5. Select **Resources**. Select **Manage Resources**. A list of resources that are available is displayed. Select your managed cluster and click **Save**. For more information, see [Add resources to a team](#).

Log in to IBM Cloud Pak for Multicloud Management.

8. Log in to IBM Cloud Pak for Multicloud Management as a user that was onboarded with the *PRIMARY\_OWNER* role and create your first namespace as part of the login process:

```
cloudctl login -a https://IP_ADDRESS:8443 -u userID -p password
```

## Online installation of IBM Cloud App Management with IBM Cloud Pak for Multicloud Management

This section includes procedures to install IBM Cloud App Management with IBM Cloud Pak for Multicloud Management online.

### Online proof-of-concept installation of IBM Cloud App Management with IBM Cloud Pak for Multicloud Management

Use this procedure to install the Cloud App Management server, and the ICAM klusterlet on an existing managed cluster. It is a proof-of-concept (POC) type installation, which installs the product quickly online to demonstrate the benefits of IBM Cloud App Management within IBM Cloud Pak for Multicloud Management.

#### Prerequisites

1. Access to two environments:
  - Environment 1 - the hub cluster: The Cloud App Management server will be installed on the hub cluster in this procedure. This environment must previously consist of the following prerequisites:
    - IBM Cloud Pak for Multicloud Management is installed.  
For more information, see the "IBM Cloud Pak for Multicloud Management Core packages" section in the IBM Cloud Pak for Multicloud Management [Passport Advantage part numbers](#) topic.
    - Red Hat OpenShift V3.11 is installed.
    - 2 compute nodes, 12 cores, 35 GB of RAM, and 2 compute nodes.
  - Environment 2 - the managed cluster: The ICAM klusterlet will be installed on a separate managed cluster in this procedure. This environment must previously consist of the following prerequisites:
    - IBM Cloud Private V3.2.1 and latest fix pack is installed.
    - Red Hat OpenShift V3.11 is installed.
    - 8 cores, 20 GB of RAM, 50 GB of spare disk space on one compute node.
2. You must have CLI access to the master nodes of your two clusters. You must install the oc or kubectl programs. The bash shell must be available.  
To install the Helm CLI, see [Installing the Helm CLI \(helm\)](#). To install the Kubernetes CLI, kubectl, see [Installing the Kubernetes CLI \(kubectl\)](#).
3. You must have access to an LDAP directory with any users and teams onboarded. For more information, see ["Onboarding LDAP users" on page 94](#).
4. The ICAM klusterlet cluster must be added as a managed cluster on the Cloud App Management server cluster. To verify that this step is complete, navigate to your /multicloud/clusters/ address on the IBM Cloud Pak for Multicloud Management dashboard.

#### Entitled Registry setup steps

You must complete the following Entitled Registry setup steps before continuing.

1. Get your IBM entitled registry API key for your Cloud App Management image from IBM Entitled Registry:
  - a. Log in to the [MyIBM Container Software Library](#) with the IBMid and password that is associated with the entitled software.
  - b. If you are not on the entitlement page, click **Get entitlement key**.
  - c. In the **Entitlement key** section, click **Copy key** to copy the entitled registry API key that is displayed.
2. Create the Docker pull secret name to pull the IBM Cloud App Management image from the IBM Cloud entitled registry:

```
kubectl create secret docker-registry secret_name --docker-username=cp \  
--docker-password=entitled_registry_api_key --docker-email=email_address --docker-  
server=cp.icr.io -n my_namespace
```

3. Patch your serviceaccount with the Docker pull secret that you created in the previous step.

```
kubectl patch serviceaccount default -p '{"imagePullSecrets": [{"name": "secret_name"}]}' -n my_namespace
```

4. Verify that your Docker image pull secret was added to your default service account:

```
kubectl describe serviceaccount default -n my_namespace
```

### Before you begin

1. You must open a terminal window with a running shell session on the master node of the Cloud App Management server cluster.
2. Ensure that you can log in to the Cloud App Management server cluster as the cluster administrator using the `oc login` command.

### The main steps

- [“1. Installing the Cloud App Management server” on page 97](#)
- [“2. Installing the ICAM kuberneteslet on the managed cluster” on page 99](#)
- [“3. Navigating to the ICAM dashboard and viewing your managed cluster” on page 99](#)

### 1. Installing the Cloud App Management server

1. Access the Helm charts in the catalog and configure the values by completing the following steps:
  - a. Log in to the console.
  - b. Click **Catalog**.
  - c. Select **Repositories**.
  - d. Select the **ibm-entitled-charts**.
  - e. Search for the `ibm-cloud-appmgmt-prod` Helm chart and select it.
  - f. Click **Configure**.
  - g. Under **Configuration**, configure the following parameters:

*Table 11. Configuration parameters*

Parameter name	Description/Commands	Example
<b>Helm release name</b>	Enter any alphanumeric string, which can include dashes.	icam
<b>Target namespace</b>	Enter kube-system.	kube-system
<b>Target cluster</b>	Enter local-cluster.	local-cluster
<b>License Accepted</b>	Select the checkbox.	

- h. Under **Quick start**, configure the following parameters:

*Table 12. Quick start parameters*

Parameter name	Description/Commands	Example
<b>Cloud Console FQDN</b>	<code>oc get routes icp-console -o=jsonpath='{.spec.host}'</code>	hostname.domain.com
<b>Cloud Console Port</b>	<code>oc get configmap ibmcloud-cluster-info -n kube-public -o=jsonpath='{.data.cluster_router_https_port}'</code>	443

Table 12. Quick start parameters (continued)

Parameter name	Description/Commands	Example
<b>Cloud Proxy FQDN</b>	oc get configmap ibmcloud-cluster-info -n kube-public -o=jsonpath='{.data.proxy_address}'	hostname.domain.com
<b>Cloud Master FQDN</b>	This is the same value as the Cloud Console FQDN.	hostname.domain.com
<b>Cloud Master Port</b>	This is the same value as the Cloud Console Port.	443
<b>Cluster administrator username</b>	The username of your cluster administrator.	admin

i. Under **All Parameters**, and configure the following parameters:

Table 13. All parameters

Parameter name	Description/Command	Example
<b>Image Repository</b>	This is the name of the entitled registry repository.	cp.icr.io/cp/app-mgmt
<b>Image Pull Secret Name</b>	This is the name of the secret that you are using as the image pull secret. Enter the secret name that you created in <a href="#">step 2</a> in the Entitled Registry Setup steps at the beginning of this topic.	<i>entitled-registry-pull-secret</i>
<b>Create TLS Certs</b>	Select the checkbox.	NA

2. Click **Install** to deploy the `ibm-cloud-appmgmt-prod` Helm chart.

**Note:** The installation can take 10 - 30 minutes.

3. When the installation is finished, in the console, click the hamburger menu in the upper left corner, and select **Helm Releases**. Select the helm release name that you provided in the previous step `icam` in this example.

4. Scroll to the notes section at the bottom of your release and complete **step 3 (OIDC registration)** of the Notes® section. Run the two **kubect1** commands that are displayed in step 3 in the Notes section.

**Note:** The **kubect1** commands are different in most cases. Do not copy them from this procedure. Instead, use the commands that are generated for you and your helm release. The first command registers with OIDC, and the second command adds two policy registrations. The following are example commands only.

```
COMMAND 1:
kubect1 exec -n kube-system -t `kubect1 get pods -l release=icam -n kube-system | grep "icam-ibm-cem-cem-users" | /
grep "Running" | head -n 1 | awk '{print $1}'` bash -- "/etc/oidc/oidc_reg.sh" "`echo $
(kubect1 get secret platform-oidc-credentials /
-o yaml -n kube-system | grep OAUTH2_CLIENT_REGISTRATION_SECRET: | awk '{print $2}'`"
Registering IBM Cloud Event Management identity ...

Checking registration...
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
   100    135     0   135     0     0    537     0  --:--:--  --:--:--  --:--:--   544

* Client does not exist...
Registering client...

Done.

COMMAND 2:
```

```
kubectl exec -n kube-system -t `kubectl get pods -l release=icam -n kube-system | grep "icam-ibm-cem-cem-users" | /
grep "Running" | head -n 1 | awk '{print $1}'` bash -- "/etc/oidc/registerServicePolicy.sh" /
"echo $(kubectl get secret icam-cem-service-secret -o yaml -n kube-system | grep cem-
service-id: | awk '{print $2}')" "`cloudctl tokens --access`"
Registering IBM Cloud Event Management identity ...
Checking service policy registration...
% Total      % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload   Total   Spent    Left   Speed
100    15  100    15    0     0      8      0  0:00:01  0:00:01  --:--:--    8
Adding first policy registration
Adding second policy registration
Done
```

The Cloud App Management server is now successfully installed. To validate that the release is working correctly, run the following command:

```
helm test releaseName --cleanup --tls
```

where *releaseName* is the helm release name you provided when you configured the helm chart through the catalog, see [Table 11 on page 97](#).

## 2. Installing the ICAM klusterlet on the managed cluster

To install the ICAM klusterlet on the managed cluster, follow the [“Installing the ICAM klusterlet on the managed cluster with Helm” on page 133](#) procedure.

## 3. Navigating to the ICAM dashboard and viewing your managed cluster

To navigate to the ICAM dashboard:

1. Log in to the console.
2. Click the hamburger menu in the upper left corner and select **Event Management**.
3. In the ICAM dashboard, on the Applications page, click **App Monitoring** to view the applications that are being monitored. To view the incidents that are being reported from the managed cluster, from the hamburger menu, under the **Monitoring** section, select **Incident**.

## Online full monitoring installation of IBM Cloud App Management with IBM Cloud Pak for Multicloud Management

Install and configure the Cloud App Management server online as part of the IBM Cloud Pak for Multicloud Management on the hub cluster in full monitoring mode.

### Before you begin

You must complete the following pre-installation steps:

1. Ensure that IBM Cloud Pak for Multicloud Management is installed on your hub cluster. For more information, see the "Install the IBM Cloud Pak for Multicloud Management section" in [Installing the IBM Cloud Pak for Multicloud Management offline](#) topic.
2. Install the Helm CLI. For instructions, see [Installing the Helm CLI \(helm\)](#).
3. Install the Kubernetes CLI, kubectl, and configure access to your cluster. For instructions, see [Installing the Kubernetes CLI \(kubectl\)](#).
4. Install the cloudctl command-line CLI. For instructions, see [Installing cloudctl](#).
5. You must onboard LDAP users. For more information, see [“Onboarding LDAP users” on page 94](#).
6. Configure the Elasticsearch **vm.max\_map\_count** parameter on all worker nodes by completing the following steps:
  - a. For Elasticsearch, you must set a kernel parameter on all worker nodes. These nodes are identified when you configure the persistent storage later in this procedure. You must set **vm.max\_map\_count** to a minimum value of 1048575. Set the parameter by using the `sysctl`

command to ensure that the change takes effect immediately. Run the following command on each worker node:

```
sysctl -w vm.max_map_count=1048575
```

- b. You must also set the **vm.max\_map\_count** parameter in the `/etc/sysctl.conf` file to ensure that the change is still in effect after the node is restarted.

```
vm.max_map_count=1048575
```

### Entitled Registry setup steps:

You must complete the following Entitled Registry steps before continuing:

1. Get your IBM entitled registry API key for your Cloud App Management image from IBM Entitled Registry:
  - a. Log in to the [MyIBM Container Software Library](#) with the IBMid and password that is associated with the entitled software.
  - b. If you are not on the entitlement page, click **Get entitlement key**.
  - c. In the **Entitlement key** section, click **Copy key** to copy the entitled registry API key that is displayed.
2. Create the Docker pull secret name to pull the IBM Cloud App Management image from the IBM Cloud entitled registry:

```
kubectl create secret docker-registry secret_name --docker-username=cp \
--docker-password=entitled_registry_api_key --docker-email=email_address --docker-
server=cp.icr.io -n my_namespace
```

3. Patch your serviceaccount with the Docker pull secret that you created in the previous step.

```
kubectl patch serviceaccount default -p '{"imagePullSecrets": [{"name": "secret_name"}]}' -n
my_namespace
```

4. Verify that your Docker image pull secret was added to your default service account:

```
kubectl describe serviceaccount default -n my_namespace
```

### Procedure

Complete the following steps as an administrator on your hub cluster on the infrastructure node of your IBM Cloud Pak for Multicloud Management environment.

1. As an administrator, log in to the management console.

```
cloudctl login -a my_cluster_URL -n kube-system --skip-ssl-validation
```

Where *my\_cluster\_URL* is the name that you defined for your cluster such as `https://cluster_address:443`. For future references to `masterIP`, use the value for *cluster\_address*. An example of *cluster\_address* is `icp-console.apps.hostname-icp-mst.domain.com`.

2. As an OpenShift administrator, log in to the OpenShift Container Platform:

```
oc login
```

3. Optional: Create the Cassandra auth secret: If you want Cassandra to use a superuser username and password of your choice, run the following command:

```
kubectl create secret generic my_release_name-cassandra-auth-secret -n kube-system --from-
literal=username=cassandraUser \
--from-literal=password=cassandraPass
```



Where *my\_release\_name* is the name that you are using for the Cloud App Management release, for example: *ibmcloudappgmt*, *cassandraUser* is the superuser name, and *cassandraPass* is the password. For example:

```
kubectl create secret generic ibmcloudappgmt-cassandra-auth-secret -n kube-system --from-literal=username=coolCassandraUser \
--from-literal=password=superSecretPassword
```

**Note:** Entering the password using the command line as shown in the previous examples might be insecure. Instead, you can add the password to a text file (ensure that there is no new line at the end of it), and use the text file in the command. For example:

```
kubectl create secret generic ibmcloudappgmt-cassandra-auth-secret -n kube-system --from-literal=username=coolCassandraUser \
--from-literal=password=$(cat cassandraPasswordfile.txt)
```

If you do not create the Cassandra auth secret yourself, a secret with a generated username and password is created for you.

4. Prepare Persistent storage: For each Cloud App Management service: Cassandra, Kafka, ZooKeeper, CouchDB, and Datalayer, persistent storage is required in a production grade environment.

If you need to access the storage script or other scripts, from the **Overview** tab, select the **SOURCE & TAR FILES** drop-down menu, and select the link to download the scripts.

For information about how to configure persistent storage, see:

- [Understanding Kubernetes storage](#)
- [Planning a storage solution](#)
- [Planning persistent storage](#)

5. Log in to the console of your target cluster.
6. Click **Catalog** in the upper right corner.
7. Select **Repositories**.
8. Select the **ibm-entitled-charts**.
9. Search for the *ibm-cloud-appgmt-prod* Helm chart and select it.
10. Click **Configure**.
11. Configure all parameters under **Configuration**.  
See [Table 1](#) for more information.
12. Under the **Parameters** area, expand **All parameters** and configure the required parameters.  
See [Table 1](#) for more information.

Parameter name	Description/Commands	Example
<b>Helm release name</b>	The Helm release name. Any alphanumeric string, which can include dashes.	ibmcloudappgmt
<b>Target namespace</b>	Select the <b>kube-system</b> namespace.	kube-system
<b>Target cluster</b>	The cluster that the <i>ibm-cloud-appgmt-prod</i> Helm chart is being installed on	
<b>License Accepted</b>	Select the checkbox.	✓

Table 14. Helm chart configuration parameters (continued)


Parameter name	Description/Commands	Example
<b>Create CRD</b>	Required for multi-cloud integrations	

Table 14. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<p><b>Create TLS Certs</b></p>	<ul style="list-style-type: none"> <li>If you are completing an online installation using Entitled Registry, select the check box.</li> <li>If you are completing an offline installation (downloading the PPA from IBM Passport Advantage) and you want to accept the default setting for creating TLS certs, which means this check box is not selected, you must run the <code>make-ca-cert-icam.sh</code> script to create the Ingress secrets. It is important that you run this script before you install the <code>ibm-cloud-appmgmt-prod</code> Helm chart. An example of how to run the <code>make-ca-cert-icam.sh</code> is: <pre data-bbox="857 1066 1138 1339"> ./ibm_cloud_pak/ pak_extensions/lib/ make-ca-cert- icam.sh cloud_Proxy_FQDN my_release_name \ kube-system icam- ingress-tls icam- ingress-client icam-ingress- artifacts </pre> <p data-bbox="857 1354 1138 1732">Where <code>cloud_Proxy_FQDN</code> is the FQDN of your IBM Cloud Pak for Multicloud Management Proxy, and <code>my_release_name</code> is the name of the Cloud App Management Helm Chart, for example: <code>ibmcloudappmgmt</code>.</p> </li> <li>If you are completing an offline installation (downloading the PPA from IBM Passport Advantage) and you want the installer to automatically create the Ingress secrets for you, then select the check box. You do not need to run the make-</li> </ul>	

Table 14. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Resource monitoring</b>	If you are installing in full monitoring mode, select the checkbox . Ensure that this check box is not checked for eventing only.	
<b>Resource Analytics</b>		
<b>Cloud Console FQDN</b>	<pre>oc get routes icp-console -o=jsonpath='{.spec.hosts}'</pre>	icp-console.hostname-icp-mst.test.com
<b>Cloud Console Port</b>	<pre>oc get configmap ibmcloud-cluster-info -n kube-public \-o=jsonpath='{.data.cluster_router_https_port}'</pre>	443
<b>Cloud Console Client Secret Name</b>	This field must be left empty.	
<b>Cloud Console TLS Secret Name</b>	This field must be left empty.	
<b>Cloud Proxy FQDN</b>	<pre>oc get configmap ibmcloud-cluster-info -n kube-public -o=jsonpath='{.data.proxy_address}'</pre>	icp-proxy.apps.hardymarmoset-icp-mst.test.com
<b>Cloud Proxy Client Secret</b>	Leave the default value.	Default value: icam-ingress-client
<b>Cloud Proxy TLS Secret</b>	Leave the default value.	Default value: icam-ingress-tls
<b>Cluster Master FQDN</b>	<pre>oc get routes icp-console -o=jsonpath='{.spec.hosts}'</pre>	icp-console.apps.hardymarmoset-icp-mst.fyre.ibm.com
<b>Cluster Master Port</b>	<pre>oc get configmap ibmcloud-cluster-info -n kube-public \-o=jsonpath='{.data.cluster_router_https_port}'</pre>	443

Table 14. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Cluster Master CA</b>	<p>Optional: If you provided your own certificate for the IBM Cloud Pak for Multicloud Management ingress, you must create a ConfigMap containing the certificate authority's certificate in PEM format (for example: <code>kubectl create configmap master-ca --from-file=./ca.pem</code>) and set this value to the name of this ConfigMap. If you did not provide your own certificate, leave this value empty.</p>	
<b>Host Alias - Cloud Proxy</b>	<p>Optional: The IP address of the IBM Cloud Pak for Multicloud Management proxy. It is used where the DNS does not resolve the IBM Cloud Pak for Multicloud Management proxy's fully qualified domain name (FQDN). It can be determined by running:</p> <pre data-bbox="833 1157 1146 1346"> kubectl get no -l   proxy=true -   o=jsonpath='{ \$.item s[*].status.address s[? (@.type=="InternalIP" )].address }'</pre>	10.21.17.70
<b>Product Deployment Size</b>	<p>Determine cluster resource requests and limits for the product. See <a href="#">“Planning hardware and sizing”</a> on page 77 for more information.</p>	test_amd64

Table 14. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Image Repository</b>	<ul style="list-style-type: none"> <li>If you completing an offline installation (downloading the PPA file from IBM Passport Advantage), for OpenShift V4.2 installations, delete the default value and replace it with <code>image-registry.openshift-image-registry.svc:5000/kube-system</code>. For OpenShift V3.11 installations, delete the default value and replace it with the shell expansion of <code>\$(oc registry info)/kube-system</code></li> <li>If you are completing an online installation using an IBM entitled registry, it is the name of the entitled registry repository.</li> </ul>	<ul style="list-style-type: none"> <li>For an offline installation: an example for OpenShift V4.2 value is: <code>image-registry.openshift-image-registry.svc:5000/kube-system</code>. An example for OpenShift V3.11 value is: <code>docker-registry.default.svc:5000/kube-system</code></li> <li>For an online installation, enter: <code>cp.icr.io/cp/app-mgmt</code></li> </ul>
<b>Image Pull Secret Name</b>	<p>If you completing an online installation using Entitled Registry, you must enter the secret name that you created in the Entitled Registry Setup steps. If you are not completing an online installation, leave this field empty.</p>	
<b>Image Prefix</b>	<p>Leave this field empty.</p>	
<b>Image pull policy</b>	<p>Leave the default value: <code>IfNotPresent</code></p>	<p><code>IfNotPresent</code></p>
<b>Cassandra Replicas</b>	<p>Configure the Cassandra Replicas.</p>	<p>1</p>

Table 14. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Default Storage Class</b>	Default storage class for the product. If the other STORAGECLASS values, for example, Cassandra Data STORAGECLASS are set to default, they use the value that is provided here. If <b>Default Storage Class</b> is left empty, the environment's default storage class is used.	glusterfs/rook-Cepheus/ or leave it empty
<b>Cassandra Data STORAGECLASS</b>	The storage class name that is used by Cassandra data PersistentVolumeClaims (PVCs). If it is set to <b>default</b> , the StorageClass name that is defined in <b>Default Storage Class</b> is used.	default
<b>Cassandra Backup STORAGECLASS</b>	The storage class name that is used by Cassandra backup PVCs. If it is set to <b>none</b> , the backup volume is disabled. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	none
<b>CouchDB Data STORAGECLASS</b>	The storage class name that is used by CouchDB PersistentVolumeClaims PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default
<b>Datalayer Jobs STORAGECLASS</b>	The storage class name that is used by Datalayer PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default
<b>Elasticsearch Data STORAGECLASS</b>	The storage class name that is used by Elasticsearch PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default

Table 14. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Kafka Data STORAGECLASS</b>	The storage class name that is used by Kafka PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default
<b>Zookeeper Data STORAGECLASS</b>	The storage class name that is used by Zookeeper PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default
...	Leave everything up to but not including the <b>Product Name</b> parameters as default values.	default
Under CEM Configuration, configure <b>Product Name</b>	<ul style="list-style-type: none"> <li>For an eventing only installation, the <b>Product Name</b> is Event Management for IBM Multicloud Manager.</li> <li>For a full monitoring installation, the <b>Product Name</b> is IBM Cloud App Management for Multicloud Manger</li> </ul>	<b>IBM Cloud App Management for Multicloud Manger</b>
...	Leave everything up to but not including the <b>Cluster administrator Username</b> parameter as default values.	Default value
<b>Cluster administrator Username</b>	The username of your cluster administrator.	admin
...	Leave all remaining parameters as default values.	Default value

13. Click **Install** to deploy the `ibm-cloud-appmgmt-prod` Helm chart.

### Results

The Cloud App Management server is successfully installed with IBM Cloud Pak for Multicloud Management in full monitoring mode. If you want to optionally verify your installation, you can run the `collectContainerLogs.sh` script, which collects the installation logs and outputs them to a



diagnostic file. For more information about running this script, see [“Collecting the server logs for IBM Support” on page 1417](#).

### What to do next

Complete the following postinstallation steps to complete the OIDC registration:

1. In the console, click the menu in the upper left corner, and select **Monitor health > Helm Releases**, then select the release name that you gave for IBM Cloud App Management.
2. Scroll to the notes section at the bottom of your release and complete **step 3 (OIDC registration)** of the Notes section. Run the two **kubect1** commands that are displayed in step 3 in the Notes section.

**Note:** The **kubect1** commands are different in most cases, do not copy and paste them from this procedure. Instead, use the commands that are generated for you and your helm release. The first command registers with OIDC, and the second command adds two policy registrations. The following are example commands only.

```
COMMAND 1:
kubect1 exec -n kube-system -t `kubect1 get pods -l release=icam -n kube-system | grep "icam-ibm-cem-cem-users" | /
grep "Running" | head -n 1 | awk '{print $1}'` bash -- "/etc/oidc/oidc_reg.sh" "`echo $(kubect1 get secret platform-oidc-credentials /
-o yaml -n kube-system | grep OAUTH2_CLIENT_REGISTRATION_SECRET: | awk '{print $2}'))`"
Registering IBM Cloud Event Management identity ...

Checking registration...
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
   Dload  Upload   Total             Spent    Left     Speed
100    135      0   135      0      0     537      0  --:--:--  --:--:--  --:--:--    544

* Client does not exist...
Registering client...

Done.

COMMAND 2:
kubect1 exec -n kube-system -t `kubect1 get pods -l release=icam -n kube-system | grep "icam-ibm-cem-cem-users" | /
grep "Running" | head -n 1 | awk '{print $1}'` bash -- "/etc/oidc/registerServicePolicy.sh" /
"`echo $(kubect1 get secret icam-cem-service-secret -o yaml -n kube-system | grep cem-
service-id: | awk '{print $2}'))" "`cloudctl tokens --access`"
Registering IBM Cloud Event Management identity ...
Checking service policy registration...
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
   Dload  Upload   Total             Spent    Left     Speed
100    15 100    15      0      0      8      0  0:00:01  0:00:01  --:--:--      8
Adding first policy registration
Adding second policy registration
Done
```

3. Now that the Cloud App Management server is successfully installed with IBM Cloud Pak for Multicloud Management, next you can deploy the ICAM klusterlet to monitor applications in your IBM Cloud Pak for Multicloud Management environment. For more information, see [“Installing the ICAM klusterlet on the managed cluster with Helm” on page 133](#) to deploy the ICAM klusterlet using its Helm chart or [“Installing the ICAM klusterlet on the managed cluster without helm” on page 135](#) for a non-Helm installation.

## Offline installation of IBM Cloud App Management with IBM Cloud Pak for Multicloud Management

This section includes procedures to install and uninstall IBM Cloud App Management with IBM Cloud Pak for Multicloud Management offline.

### Offline eventing only installation of IBM Cloud App Management with IBM Cloud Pak for Multicloud Management

You can install and configure the Cloud App Management server with IBM Cloud Pak for Multicloud Management on the hub cluster for event management only. If you want to, you can update this

configuration from eventing only to full monitoring mode. Access the instructions to complete this update at the end of this topic.

## Before you begin

You must complete the following pre-installation steps:

1. Ensure that IBM Cloud Pak for Multicloud Management is installed on your hub cluster. For more information, see the "Install the IBM Cloud Pak for Multicloud Management section" in [Installing the IBM Cloud Pak for Multicloud Management offline](#) topic.
2. Install the Helm CLI. For instructions, see [Installing the Helm CLI \(helm\)](#).
3. Install the Kubernetes CLI, kubectl, and configure access to your cluster. For instructions, see [Installing the Kubernetes CLI \(kubectl\)](#).
4. Install the cloudctl command-line CLI. For instructions, see [Installing cloudctl](#).
5. You must onboard LDAP users. For more information, see ["Onboarding LDAP users" on page 94](#).
6. If you are using Red Hat OpenShift V4.2, the docker CLI might not be installed. To check whether it is installed, run the following command:

```
which docker; echo $?
```

If the result returned is 0, the docker CLI is installed. If the docker CLI is not installed, complete the following steps to install it:

- a. Download the docker CLI installation file from [IBM Passport Advantage®](#) to the infrastructure node. To check which node is the infrastructure node, run;

```
kubectl get nodes
```

Search for the installation file by using its part number: CC3KUEN. Its file name is `icp-docker-18.09.7_x86_64.bin`.

- b. Run the following command to make the binary executable:

```
chmod +x icp-docker-18.09.7_x86_64.bin
```

- c. Run the binary.

7. Configure the Elasticsearch **vm.max\_map\_count** parameter on all worker nodes by completing the following steps:

- a. For Elasticsearch, you must set a kernel parameter on all worker nodes. These nodes are identified when you configure the persistent storage later in this procedure. You must set **vm.max\_map\_count** to a minimum value of 1048575. Set the parameter by using the `sysctl` command to ensure that the change takes effect immediately. Run the following command on each worker node:

```
sysctl -w vm.max_map_count=1048575
```

- b. You must also set the **vm.max\_map\_count** parameter in the `/etc/sysctl.conf` file to ensure that the change is still in effect after the node is restarted.

```
vm.max_map_count=1048575
```

## Procedure

Complete the following steps as an administrator on your hub cluster on the infrastructure node of your IBM Cloud Pak for Multicloud Management environment.

1. Create a directory, which you can use to save the installation files to and complete installation steps. For this procedure, an example directory that is called `install_dir` is used.

```
mkdir install_dir/
```

2. Locate the Cloud App Management server eventing only installation image file on [IBM Passport Advantage](#). Download the installation image file to the `install_dir` directory. Depending on your platform, choose one of the following images on [IBM Passport Advantage](#). Find the installation image by searching for it using its part number.

- For Linux on Power, choose the Multicluster Event Management Server on PlinuxLE (part number: CC4KYEN): `icam_ppa_2019.4.0_prod_lite_ppc64.tar.gz`
- For Linux x86\_64, choose the Multicluster Event Management Server on AMD64 (part number: CC4L0EN) `icam_ppa_2019.4.0_prod_lite_amd64.tar.gz`

For more information about the IBM Cloud App Management components in the context of IBM Cloud Pak for Multicloud Management, see the IBM Cloud App Management for IBM Cloud Pak for Multicloud Management packages section in the [Passport Advantage part numbers](#) topic.

**Note:** The Cloud App Management server with IBM Cloud Pak for Multicloud Management must be installed in the `kube-system` namespace. Throughout this procedure, if prompted for a namespace, you must use `kube-system`.

3. As an administrator, log in to the management console.

```
cloudctl login -a my_cluster_URL -n kube-system --skip-ssl-validation
```

Where `my_cluster_URL` is the name that you defined for your cluster such as `https://cluster_address:443`. For future references to `masterIP`, use the value for `cluster_address`. An example of `cluster_address` is `icp-console.apps.hostname-icp-mst.domain.com`.

4. As an OpenShift administrator, log in to the OpenShift Container Platform:

```
oc login
```

5. Log in to the Docker registry:

```
docker login $(oc registry info) -u $(oc whoami) -p $(oc whoami -t)
```

**Note:** For OpenShift V4.2 or later environment, if you get a `x509: certificate signed by unknown authority` error, you must complete the step 6 to 10 also. Ignore these steps for OpenShift V3.11.

6. Enable the default route:

```
oc patch configs.imageregistry.operator.openshift.io/cluster --patch '{"spec":{"defaultRoute":true}}' --type=merge
```

7. Obtain the default route for the registry:

```
oc get route default-route -n openshift-image-registry --template='{{ .spec.host }}'
```

8. Create and edit the `/etc/docker/daemon.json` file to include the domain:

```
{
  "insecure-registries": ["DEFAULT_ROUTE"]
}
```

9. Restart the docker daemon:

```
systemctl restart docker
```

10. Log in using the route and a user without a colon in the name:

```
docker login $DEFAULT_ROUTE -u $USER -p $(oc whoami -t)
```

11. Load the Passport Advantage Archive (PPA) file installation image file into IBM's Docker registry:

```
cloudctl catalog load-archive --archive ./installation_image_file --registry $(oc registry info)/kube-system
```

where `installation_image_file` is the compressed Cloud App Management server eventing only installation image file that you downloaded in step 2 of this procedure.

12. Extract the Helm charts from the Passport Advantage Archive (PPA) file into the `install_dir` directory. The

```
cd install_dir
tar -xvf ../installation_image_file charts
tar -xvf ../charts/ibm-cloud-appmgmt-prod-1.6.0.tgz
```

The `charts` value is required to ensure the `tar` command extracts only the charts directory from the `installation_image_file` file. Otherwise, all the images are extracted, which might cause space issues.

13. Change directory to the `ibm-cloud-appmgmt-prod` directory.

```
cd install_dir/ibm-cloud-appmgmt-prod
```

14. Optional: Create the Cassandra auth secret: If you want Cassandra to use a superuser username and password of your choice, run the following command:

```
kubectl create secret generic my_release_name-cassandra-auth-secret -n kube-system --from-literal=username=cassandraUser \
--from-literal=password=cassandraPass
```

Where `my_release_name` is the name that you are using for the Cloud App Management release, for example: `ibmcloudappmgmt`, `cassandraUser` is the superuser name, and `cassandraPass` is the password. For example:

```
kubectl create secret generic ibmcloudappmgmt-cassandra-auth-secret -n kube-system --from-literal=username=coolCassandraUser \
--from-literal=password=superSecretPassword
```

**Note:** Entering the password using the command line as shown in the previous examples might be insecure. Instead, you can add the password to a text file (ensure that there is no new line at the end of it), and use the text file in the command. For example:

```
kubectl create secret generic ibmcloudappmgmt-cassandra-auth-secret -n kube-system --from-literal=username=coolCassandraUser \
--from-literal=password=$(cat cassandraPasswordfile.txt)
```

If you do not create the Cassandra auth secret yourself, a secret with a generated username and password is created for you.

15. Prepare Persistent storage: Persistent storage is required in a production environment. IBM highly recommends local storage, that is, persistent volumes backed by local disks or partitions. For information about how to configure persistent storage, see [Understanding Kubernetes storage](#). Cloud App Management includes an optional `prepare-pv.sh` script that you can use to create storage classes and persistent volumes backed by local storage. To use this script, complete the following steps:

- a. Locate the script in the Helm chart directory. It is in `install_dir/ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh`.
- b. Execute the script without any parameters to see the usage instructions:

```
./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh
```

- c. Identify the correct parameter values and run the script with the parameter set. For example:

```
./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh --size1_amd64 --
releasename ibmcloudappmgmt \
--cassandraNode 10.10.10.1 --zookeeperNode 10.10.10.2 --kafkaNode 10.10.10.3 --
couchdbNode 10.10.10.4 --datalayerNode \
10.10.10.5 --elasticSearchNode 10.10.10.6 --local
```

- d. Create the persistent volumes and storage classes from the resource files:

```
kubectl create -f ibm-cloud-appmgmt-prod/ibm_cloud_pak/yaml/
```

- e. When you are configuring the storage class parameters later in this procedure (see [Table 15 on page 113](#) table for more information), you must provide the storage classes that you defined for

each of the different stateful sets. For example, if you used `--CassandraClass myCassandraStorageClass`, you must provide `myCassandraStorageClass` as the value for `Cassandra Storage Class`. If you did not specify custom storage classes and you used the defaults, you must provide the default values that the `prepare-pv.sh` script uses: `my_release_name-local-storage-elasticsearch`, `my_release_name-local-storage-kafka`, and so on. To verify which storage classes are created and are available to use, run **kubect1 get storageclass**.

16. Log in to the console of your target cluster.
17. Click **Catalog** in the upper right corner.
18. Search for the `ibm-cloud-appmgmt-prod` Helm chart and select it.
19. Click **Configure**.
20. Configure all parameters under **Configuration**.  
See [Table 1](#) for more information.
21. Under the **Parameters** area, expand **All parameters** and configure the required parameters.  
See [Table 1](#) for more information.

<i>Table 15. Helm chart configuration parameters</i>		
<b>Parameter name</b>	<b>Description/Commands</b>	<b>Example</b>
<b>Helm release name</b>	The Helm release name. Any alphanumeric string, which can include dashes.	ibmcloudappmgmt
<b>Target namespace</b>	Select the <b>kube-system</b> namespace.	kube-system
<b>Target cluster</b>	The cluster that the <code>ibm-cloud-appmgmt-prod</code> Helm chart is being installed on	
<b>License Accepted</b>	Select the checkbox.	✓
<b>Create CRD</b>	Required for multi-cloud integrations	✓

Table 15. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<p><b>Create TLS Certs</b></p>	<ul style="list-style-type: none"> <li>If you are completing an online installation using Entitled Registry, select the check box.</li> <li>If you are completing an offline installation (downloading the PPA from IBM Passport Advantage) and you want to accept the default setting for creating TLS certs, which means this check box is not selected, you must run the <code>make-ca-cert-icam.sh</code> script to create the Ingress secrets. It is important that you run this script before you install the <code>ibm-cloud-appmgmt-prod</code> Helm chart. An example of how to run the <code>make-ca-cert-icam.sh</code> is:           <pre data-bbox="857 1066 1138 1339"> ./ibm_cloud_pak/ pak_extensions/lib/ make-ca-cert- icam.sh cloud_Proxy_FQDN my_release_name \ kube-system icam- ingress-tls icam- ingress-client icam-ingress- artifacts           </pre> <p data-bbox="857 1360 1138 1738">Where <code>cloud_Proxy_FQDN</code> is the FQDN of your IBM Cloud Pak for Multicloud Management Proxy, and <code>my_release_name</code> is the name of the Cloud App Management Helm Chart, for example: <code>ibmcloudappmgmt</code>.</p> </li> <li>If you are completing an offline installation (downloading the PPA from IBM Passport Advantage) and you want the installer to automatically create the Ingress secrets for you, then select the check box. You do not need to run the <code>make-</code></li> </ul>	

Table 15. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Resource monitoring</b>	If you are installing in full monitoring mode, select the checkbox . Ensure that this check box is not checked for eventing only.	
<b>Resource Analytics</b>		
<b>Cloud Console FQDN</b>	<pre>oc get routes icp-console -o=jsonpath='{.spec.hosts}'</pre>	icp-console.hostname-icp-mst.test.com
<b>Cloud Console Port</b>	<pre>oc get configmap ibmcloud-cluster-info -n kube-public \-o=jsonpath='{.data.cluster_router_https_port}'</pre>	443
<b>Cloud Console Client Secret Name</b>	This field must be left empty.	
<b>Cloud Console TLS Secret Name</b>	This field must be left empty.	
<b>Cloud Proxy FQDN</b>	<pre>oc get configmap ibmcloud-cluster-info -n kube-public -o=jsonpath='{.data.proxy_address}'</pre>	icp-proxy.apps.hardymarmoset-icp-mst.test.com
<b>Cloud Proxy Client Secret</b>	Leave the default value.	Default value: icam-ingress-client
<b>Cloud Proxy TLS Secret</b>	Leave the default value.	Default value: icam-ingress-tls
<b>Cluster Master FQDN</b>	<pre>oc get routes icp-console -o=jsonpath='{.spec.hosts}'</pre>	icp-console.apps.hardymarmoset-icp-mst.fyre.ibm.com
<b>Cluster Master Port</b>	<pre>oc get configmap ibmcloud-cluster-info -n kube-public \-o=jsonpath='{.data.cluster_router_https_port}'</pre>	443

Table 15. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Cluster Master CA</b>	<p>Optional: If you provided your own certificate for the IBM Cloud Pak for Multicloud Management ingress, you must create a ConfigMap containing the certificate authority's certificate in PEM format (for example: <code>kubectl create configmap master-ca --from-file=./ca.pem</code>) and set this value to the name of this ConfigMap. If you did not provide your own certificate, leave this value empty.</p>	
<b>Host Alias - Cloud Proxy</b>	<p>Optional: The IP address of the IBM Cloud Pak for Multicloud Management proxy. It is used where the DNS does not resolve the IBM Cloud Pak for Multicloud Management proxy's fully qualified domain name (FQDN). It can be determined by running:</p> <pre data-bbox="833 1157 1146 1346"> kubectl get no -l   proxy=true -   o=jsonpath='{ \$.item s[*].status.address s[? (@.type=="InternalIP" )].address }'</pre>	10.21.17.70
<b>Product Deployment Size</b>	<p>Determine cluster resource requests and limits for the product. See <a href="#">“Planning hardware and sizing”</a> on page 77 for more information.</p>	test_amd64



Table 15. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Image Repository</b>	<ul style="list-style-type: none"> <li>If you completing an offline installation (downloading the PPA file from IBM Passport Advantage), for OpenShift V4.2 installations, delete the default value and replace it with <code>image-registry.openshift-image-registry.svc:5000/kube-system</code>. For OpenShift V3.11 installations, delete the default value and replace it with the shell expansion of <code>\$(oc registry info)/kube-system</code></li> <li>If you are completing an online installation using an IBM entitled registry, it is the name of the entitled registry repository.</li> </ul>	<ul style="list-style-type: none"> <li>For an offline installation: an example for OpenShift V4.2 value is: <code>image-registry.openshift-image-registry.svc:5000/kube-system</code>. An example for OpenShift V3.11 value is: <code>docker-registry.default.svc:5000/kube-system</code></li> <li>For an online installation, enter: <code>cp.icr.io/cp/app-mgmt</code></li> </ul>
<b>Image Pull Secret Name</b>	<p>If you completing an online installation using Entitled Registry, you must enter the secret name that you created in the Entitled Registry Setup steps. If you are not completing an online installation, leave this field empty.</p>	
<b>Image Prefix</b>	<p>Leave this field empty.</p>	
<b>Image pull policy</b>	<p>Leave the default value: <code>IfNotPresent</code></p>	<p><code>IfNotPresent</code></p>
<b>Cassandra Replicas</b>	<p>Configure the Cassandra Replicas.</p>	<p>1</p>

Table 15. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Default Storage Class</b>	Default storage class for the product. If the other STORAGECLASS values, for example, Cassandra Data STORAGECLASS are set to default, they use the value that is provided here. If <b>Default Storage Class</b> is left empty, the environment's default storage class is used.	glusterfs/rook-Cepheus/ or leave it empty
<b>Cassandra Data STORAGECLASS</b>	The storage class name that is used by Cassandra data PersistentVolumeClaims (PVCs). If it is set to <b>default</b> , the StorageClass name that is defined in <b>Default Storage Class</b> is used.	default
<b>Cassandra Backup STORAGECLASS</b>	The storage class name that is used by Cassandra backup PVCs. If it is set to <b>none</b> , the backup volume is disabled. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	none
<b>CouchDB Data STORAGECLASS</b>	The storage class name that is used by CouchDB PersistentVolumeClaims PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default
<b>Datalayer Jobs STORAGECLASS</b>	The storage class name that is used by Datalayer PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default
<b>Elasticsearch Data STORAGECLASS</b>	The storage class name that is used by Elasticsearch PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default

<i>Table 15. Helm chart configuration parameters (continued)</i>		
<b>Parameter name</b>	<b>Description/Commands</b>	<b>Example</b>
<b>Kafka Data STORAGECLASS</b>	The storage class name that is used by Kafka PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default
<b>Zookeeper Data STORAGECLASS</b>	The storage class name that is used by Zookeeper PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default
...	Leave everything up to but not including the <b>Product Name</b> parameters as default values.	default
Under CEM Configuration, configure <b>Product Name</b>	<ul style="list-style-type: none"> <li>For an eventing only installation, the <b>Product Name</b> is Event Management for IBM Multicloud Manager.</li> <li>For a full monitoring installation, the <b>Product Name</b> is IBM Cloud App Management for Multicloud Manger</li> </ul>	<b>IBM Cloud App Management for Multicloud Manger</b>
...	Leave everything up to but not including the <b>Cluster administrator Username</b> parameter as default values.	Default value
<b>Cluster administrator Username</b>	The username of your cluster administrator.	admin
...	Leave all remaining parameters as default values.	Default value

22. Click **Install** to deploy the `ibm-cloud-appmgmt-prod` Helm chart.

### Results

After approximately 20 -30 minutes, the Cloud App Management server is successfully installed with IBM Cloud Pak for Multicloud Management for eventing only. If you want to optionally verify your installation, you can run the `collectContainerLogs.sh` script, which collects the installation logs and outputs

them to a diagnostic file. For more information about running this script, see [“Collecting the server logs for IBM Support”](#) on page 1417.

If you want to, you can update from eventing only to full monitoring mode by completing the steps in [Updating from Eventing only to full monitoring mode](#) next.

## What to do next

Complete the following postinstallation steps:

1. In the console, click the menu in the upper left corner, and select **Monitor health > Helm Releases**, then select the release name that you gave for IBM Cloud App Management.
2. Scroll to the notes section at the bottom of your release and complete **step 3 (OIDC registration)** of the Notes section. Run the two **kubect1** commands that are displayed in step 3 in the Notes section.

**Note:** The **kubect1** commands are different in most cases, do not copy and paste them from this procedure. Instead, use the commands that are generated for you and your helm release. The first command registers with OIDC, and the second command adds two policy registrations. The following are example commands only.

```
COMMAND 1:
kubect1 exec -n kube-system -t `kubect1 get pods -l release=icam -n kube-system | grep "icam-ibm-cem-cem-users" | /
grep "Running" | head -n 1 | awk '{print $1}'` bash -- "/etc/oidc/oidc_reg.sh" "`echo $(kubect1 get secret platform-oidc-credentials /
-o yaml -n kube-system | grep OAUTH2_CLIENT_REGISTRATION_SECRET: | awk '{print $2}'))`"
Registering IBM Cloud Event Management identity ...

Checking registration...
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           0         0     0         0          0      0      0     0
100    135     0    135     0     0     537     0  ---:---:--  ---:---:--  ---:---:--    544

* Client does not exist...
Registering client...

Done.

COMMAND 2:
kubect1 exec -n kube-system -t `kubect1 get pods -l release=icam -n kube-system | grep "icam-ibm-cem-cem-users" | /
grep "Running" | head -n 1 | awk '{print $1}'` bash -- "/etc/oidc/registerServicePolicy.sh" /
"`echo $(kubect1 get secret icam-cem-service-secret -o yaml -n kube-system | grep cem-
service-id: | awk '{print $2}'))`" "`cloudctl tokens --access`"
Registering IBM Cloud Event Management identity ...
Checking service policy registration...
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           0         0     0         0          0      0      0     0
100    15 100    15     0     0     8     0  0:00:01  0:00:01  ---:---:--     8
Adding first policy registration
Adding second policy registration
Done
```

## Offline full monitoring installation of IBM Cloud App Management with IBM Cloud Pak for Multicloud Management

Install and configure the Cloud App Management server as part of the IBM Cloud Pak for Multicloud Management on the hub cluster in full monitoring mode.

### Before you begin

You must complete the following pre-installation steps:

1. Ensure that IBM Cloud Pak for Multicloud Management is installed on your hub cluster. For more information, see the "Install the IBM Cloud Pak for Multicloud Management section" in [Installing the IBM Cloud Pak for Multicloud Management offline](#) topic.
2. Install the Helm CLI. For instructions, see [Installing the Helm CLI \(helm\)](#).
3. Install the Kubernetes CLI, kubect1, and configure access to your cluster. For instructions, see [Installing the Kubernetes CLI \(kubect1\)](#).

4. Install the `cloudctl` command-line CLI. For instructions, see [Installing cloudctl](#).
5. You must onboard LDAP users. For more information, see [“Onboarding LDAP users”](#) on page 94.
6. If you are using Red Hat OpenShift V4.2, the `docker` CLI might not be installed. To check whether it is installed, run the following command:

```
which docker; echo $?
```

If the result returned is `0`, the `docker` CLI is installed. If the `docker` CLI is not installed, complete the following steps to install it:

- a. Download the `docker` CLI installation file from [IBM Passport Advantage®](#) to the infrastructure node. To check which node is the infrastructure node, run;

```
kubectl get nodes
```

Search for the installation file by using its part number: CC3KUEN. Its file name is `icp-docker-18.09.7_x86_64.bin`.

- b. Run the following command to make the binary executable:

```
chmod +x icp-docker-18.09.7_x86_64.bin
```

- c. Run the binary.

7. Configure the Elasticsearch `vm.max_map_count` parameter on all worker nodes by completing the following steps:

- a. For Elasticsearch, you must set a kernel parameter on all worker nodes. These nodes are identified when you configure the persistent storage later in this procedure. You must set `vm.max_map_count` to a minimum value of 1048575. Set the parameter by using the `sysctl` command to ensure that the change takes effect immediately. Run the following command on each worker node:

```
sysctl -w vm.max_map_count=1048575
```

- b. You must also set the `vm.max_map_count` parameter in the `/etc/sysctl.conf` file to ensure that the change is still in effect after the node is restarted.

```
vm.max_map_count=1048575
```

## Procedure

Complete the following steps as an administrator on your hub cluster on the infrastructure node of your IBM Cloud Pak for Multicloud Management environment.

1. Create a directory, which you can use to save the installation files to and complete installation steps. For this procedure, an example directory that is called `install_dir` is used.

```
mkdir install_dir/
```

2. Locate the Cloud App Management server full monitoring installation image file on [IBM Passport Advantage®](#). Download the installation image file to the `install_dir` directory. Depending on your platform, choose one of the following images. Find the installation image by searching for it using its part number.
  - For Linux `x86_64`, choose `icam_ppa_2019.4.0_prod.tar.gz` (Part number: CC4KNEN)
  - For Linux on Power, choose `icam_ppa_2019.4.0_prod_ppc64le.tar.gz` (Part number: CC4LHEN)

For more information about the IBM Cloud App Management components in the context of IBM Cloud Pak for Multicloud Management, see the [IBM Cloud App Management for IBM Cloud Pak for Multicloud Management packages](#) section in the [Passport Advantage part numbers](#) topic.

**Note:** The Cloud App Management server with IBM Cloud Pak for Multicloud Management must be installed in the `kube-system` namespace. Throughout this procedure, if prompted for a namespace, you must use `kube-system`.

3. As an administrator, log in to the management console.

```
cloudctl login -a my_cluster_URL -n kube-system --skip-ssl-validation
```

Where *my\_cluster\_URL* is the name that you defined for your cluster such as `https://cluster_address:443`. For future references to `masterIP`, use the value for *cluster\_address*. An example of *cluster\_address* is `icp-console.apps.hostname-icp-mst.domain.com`.

4. As an OpenShift administrator, log in to the OpenShift Container Platform:

```
oc login
```

5. Log in to the Docker registry:

```
docker login $(oc registry info) -u $(oc whoami) -p $(oc whoami -t)
```

**Note:** For OpenShift V4.2 or later environment, if you get a `x509: certificate signed by unknown authority` error, you must complete the step 6 to 10 also. Ignore these steps for OpenShift V3.11.

6. Enable the default route:

```
oc patch configs.imageregistry.operator.openshift.io/cluster --patch '{"spec":{"defaultRoute":true}}' --type=merge
```

7. Obtain the default route for the registry:

```
oc get route default-route -n openshift-image-registry --template='{{ .spec.host }}'
```

8. Create and edit the `/etc/docker/daemon.json` file to include the domain:

```
{
  "insecure-registries": ["DEFAULT_ROUTE"]
}
```

9. Restart the docker daemon:

```
systemctl restart docker
```

10. Log in using the route and a user without a colon in the name:

```
docker login $DEFAULT_ROUTE -u $USER -p $(oc whoami -t)
```

11. Load the Passport Advantage Archive (PPA) file installation image file into IBM's Docker registry:

```
cloudctl catalog load-archive --archive installation_image_file --registry $(oc registry info)/kube-system
```

where *installation\_image\_file* is the compressed Cloud App Management server full monitoring installation image file that you downloaded in step 2 of this procedure.

12. Extract the Helm charts from the Passport Advantage Archive (PPA) file into the `install_dir` directory. The

```
cd install_dir
tar -xvf installation_image_file charts
tar -xvf charts/ibm-cloud-appmgmt-prod-1.6.0.tgz
```

The `charts` value is required to ensure the `tar` command extracts only the charts directory from the *installation\_image\_file* file. Otherwise, all the images are extracted, which might cause space issues.

13. Change directory to the `ibm-cloud-appmgmt-prod` directory.

```
cd install_dir/ibm-cloud-appmgmt-prod
```

14. Create the Cloud App Management ingress TLS and client secrets:

```
./ibm_cloud_pak/pak_extensions/lib/make-ca-cert-icam.sh cloud_Proxy_FQDN my_release_name \
kube-system icam-ingress-tls icam-ingress-client icam-ingress-artifacts
```

Where *cloud\_Proxy\_FQDN* is the FQDN of your IBM Cloud Pak for Multicloud Management Proxy, and *my\_release\_name* is the name of the Cloud App Management Helm Chart, for example: *ibmcloudappmgmt*.

15. Optional: Create the Cassandra auth secret: If you want Cassandra to use a superuser username and password of your choice, run the following command:

```
kubect1 create secret generic my_release_name-cassandra-auth-secret -n kube-system --from-
literal=username=cassandraUser \
--from-literal=password=cassandraPass
```

Where *my\_release\_name* is the name that you are using for the Cloud App Management release, for example: *ibmcloudappmgmt*, *cassandraUser* is the superuser name, and *cassandraPass* is the password. For example:

```
kubect1 create secret generic ibmcloudappmgmt-cassandra-auth-secret -n kube-system --from-
literal=username=coolCassandraUser \
--from-literal=password=superSecretPassword
```

**Note:** Entering the password using the command line as shown in the previous examples might be insecure. Instead, you can add the password to a text file (ensure that there is no new line at the end of it), and use the text file in the command. For example:

```
kubect1 create secret generic ibmcloudappmgmt-cassandra-auth-secret -n kube-system --from-
literal=username=coolCassandraUser \
--from-literal=password=$(cat cassandraPasswordfile.txt)
```

If you do not create the Cassandra auth secret yourself, a secret with a generated username and password is created for you.

16. Prepare Persistent storage: Persistent storage is required in a production environment. IBM highly recommends local storage, that is, persistent volumes backed by local disks or partitions. For information about how to configure persistent storage, see [Understanding Kubernetes storage](#). Cloud App Management includes an optional `prepare-pv.sh` script that you can use to create storage classes and persistent volumes backed by local storage. To use this script, complete the following steps:

- a. Locate the script in the Helm chart directory. It is in `install_dir/ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh`.
- b. Execute the script without any parameters to see the usage instructions:

```
./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh
```

- c. Identify the correct parameter values and run the script with the parameter set. For example:

```
./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh --size1_amd64 --
releasename ibmcloudappmgmt \
--cassandraNode 10.10.10.1 --zookeeperNode 10.10.10.2 --kafkaNode 10.10.10.3 --
couchdbNode 10.10.10.4 --datalayerNode \
10.10.10.5 --elasticSearchNode 10.10.10.6 --local
```

- d. Create the persistent volumes and storage classes from the resource files:

```
kubect1 create -f ibm-cloud-appmgmt-prod/ibm_cloud_pak/yaml/
```

- e. When you are configuring the storage class parameters later in this procedure (see [Table 16 on page 124](#) table for more information), you must provide the storage classes that you defined for each of the different stateful sets. For example, if you used `--CassandraClass myCassandraStorageClass`, you must provide `myCassandraStorageClass` as the value for `Cassandra Storage Class`. If you did not specify custom storage classes and you used the defaults, you must provide the default values that the `prepare-pv.sh` script uses:

`my_release_name-local-storage-elasticsearch`, `my_release_name-local-storage-kafka`, and so on. To verify which storage classes are created and are available to use, run **kubectl get storageclass**.

17. Log in to the console of your target cluster.
18. Click **Catalog** in the upper right corner.
19. Search for the `ibm-cloud-appmgmt-prod` Helm chart and select it.
20. Click **Configure**.
21. Configure all parameters under **Configuration**.  
See [Table 1](#) for more information.
22. Under the **Parameters** area, expand **All parameters** and configure the required parameters.  
See [Table 1](#) for more information.

<i>Table 16. Helm chart configuration parameters</i>		
<b>Parameter name</b>	<b>Description/Commands</b>	<b>Example</b>
<b>Helm release name</b>	The Helm release name. Any alphanumeric string, which can include dashes.	ibmcloudappmgmt
<b>Target namespace</b>	Select the <b>kube-system</b> namespace.	kube-system
<b>Target cluster</b>	The cluster that the <code>ibm-cloud-appmgmt-prod</code> Helm chart is being installed on	
<b>License Accepted</b>	Select the checkbox.	✓
<b>Create CRD</b>	Required for multi-cloud integrations	✓



Table 16. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<p><b>Create TLS Certs</b></p>	<ul style="list-style-type: none"> <li>If you are completing an online installation using Entitled Registry, select the check box.</li> <li>If you are completing an offline installation (downloading the PPA from IBM Passport Advantage) and you want to accept the default setting for creating TLS certs, which means this check box is not selected, you must run the <code>make-ca-cert-icam.sh</code> script to create the Ingress secrets. It is important that you run this script before you install the <code>ibm-cloud-appmgmt-prod</code> Helm chart. An example of how to run the <code>make-ca-cert-icam.sh</code> is:           <pre data-bbox="857 1066 1141 1339"> ./ibm_cloud_pak/ pak_extensions/lib/ make-ca-cert- icam.sh cloud_Proxy_FQDN my_release_name \ kube-system icam- ingress-tls icam- ingress-client icam-ingress- artifacts </pre> <p>Where <code>cloud_Proxy_FQDN</code> is the FQDN of your IBM Cloud Pak for Multicloud Management Proxy, and <code>my_release_name</code> is the name of the Cloud App Management Helm Chart, for example: <code>ibmcloudappmgmt</code>.</p> </li> <li>If you are completing an offline installation (downloading the PPA from IBM Passport Advantage) and you want the installer to automatically create the Ingress secrets for you, then select the check box. You do not need to run the make-</li> </ul>	

Table 16. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Resource monitoring</b>	If you are installing in full monitoring mode, select the checkbox . Ensure that this check box is not checked for eventing only.	
<b>Resource Analytics</b>		
<b>Cloud Console FQDN</b>	<pre>oc get routes icp-console -o=jsonpath='{.spec.hosts}'</pre>	icp-console.hostname-icp-mst.test.com
<b>Cloud Console Port</b>	<pre>oc get configmap ibmcloud-cluster-info -n kube-public \-o=jsonpath='{.data.cluster_router_https_port}'</pre>	443
<b>Cloud Console Client Secret Name</b>	This field must be left empty.	
<b>Cloud Console TLS Secret Name</b>	This field must be left empty.	
<b>Cloud Proxy FQDN</b>	<pre>oc get configmap ibmcloud-cluster-info -n kube-public -o=jsonpath='{.data.proxy_address}'</pre>	icp-proxy.apps.hardymarmoset-icp-mst.test.com
<b>Cloud Proxy Client Secret</b>	Leave the default value.	Default value: icam-ingress-client
<b>Cloud Proxy TLS Secret</b>	Leave the default value.	Default value: icam-ingress-tls
<b>Cluster Master FQDN</b>	<pre>oc get routes icp-console -o=jsonpath='{.spec.hosts}'</pre>	icp-console.apps.hardymarmoset-icp-mst.fyre.ibm.com
<b>Cluster Master Port</b>	<pre>oc get configmap ibmcloud-cluster-info -n kube-public \-o=jsonpath='{.data.cluster_router_https_port}'</pre>	443

Table 16. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Cluster Master CA</b>	<p>Optional: If you provided your own certificate for the IBM Cloud Pak for Multicloud Management ingress, you must create a ConfigMap containing the certificate authority's certificate in PEM format (for example: <code>kubectl create configmap master-ca --from-file=./ca.pem</code>) and set this value to the name of this ConfigMap. If you did not provide your own certificate, leave this value empty.</p>	
<b>Host Alias - Cloud Proxy</b>	<p>Optional: The IP address of the IBM Cloud Pak for Multicloud Management proxy. It is used where the DNS does not resolve the IBM Cloud Pak for Multicloud Management proxy's fully qualified domain name (FQDN). It can be determined by running:</p> <pre data-bbox="833 1157 1146 1350"> kubectl get no -l proxy=true -o=jsonpath='{ \$.items[*].status.addresses[?(@.type=="InternalIP")].address }'</pre>	10.21.17.70
<b>Product Deployment Size</b>	<p>Determine cluster resource requests and limits for the product. See <a href="#">“Planning hardware and sizing”</a> on page 77 for more information.</p>	test_amd64

Table 16. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Image Repository</b>	<ul style="list-style-type: none"> <li>If you completing an offline installation (downloading the PPA file from IBM Passport Advantage), for OpenShift V4.2 installations, delete the default value and replace it with <code>image-registry.openshift-image-registry.svc:5000/kube-system</code>. For OpenShift V3.11 installations, delete the default value and replace it with the shell expansion of <code>\$(oc registry info)/kube-system</code></li> <li>If you are completing an online installation using an IBM entitled registry, it is the name of the entitled registry repository.</li> </ul>	<ul style="list-style-type: none"> <li>For an offline installation: an example for OpenShift V4.2 value is: <code>image-registry.openshift-image-registry.svc:5000/kube-system</code>. An example for OpenShift V3.11 value is: <code>docker-registry.default.svc:5000/kube-system</code></li> <li>For an online installation, enter: <code>cp.icr.io/cp/app-mgmt</code></li> </ul>
<b>Image Pull Secret Name</b>	<p>If you completing an online installation using Entitled Registry, you must enter the secret name that you created in the Entitled Registry Setup steps. If you are not completing an online installation, leave this field empty.</p>	
<b>Image Prefix</b>	<p>Leave this field empty.</p>	
<b>Image pull policy</b>	<p>Leave the default value: <code>IfNotPresent</code></p>	<p><code>IfNotPresent</code></p>
<b>Cassandra Replicas</b>	<p>Configure the Cassandra Replicas.</p>	<p>1</p>

Table 16. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Default Storage Class</b>	Default storage class for the product. If the other STORAGECLASS values, for example, Cassandra Data STORAGECLASS are set to default, they use the value that is provided here. If <b>Default Storage Class</b> is left empty, the environment's default storage class is used.	glusterfs/rook-Cepheus/ or leave it empty
<b>Cassandra Data STORAGECLASS</b>	The storage class name that is used by Cassandra data PersistentVolumeClaims (PVCs). If it is set to <b>default</b> , the StorageClass name that is defined in <b>Default Storage Class</b> is used.	default
<b>Cassandra Backup STORAGECLASS</b>	The storage class name that is used by Cassandra backup PVCs. If it is set to <b>none</b> , the backup volume is disabled. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	none
<b>CouchDB Data STORAGECLASS</b>	The storage class name that is used by CouchDB PersistentVolumeClaims PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default
<b>Datalayer Jobs STORAGECLASS</b>	The storage class name that is used by Datalayer PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default
<b>Elasticsearch Data STORAGECLASS</b>	The storage class name that is used by Elasticsearch PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default

Table 16. Helm chart configuration parameters (continued)

Parameter name	Description/Commands	Example
<b>Kafka Data STORAGECLASS</b>	The storage class name that is used by Kafka PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default
<b>Zookeeper Data STORAGECLASS</b>	The storage class name that is used by Zookeeper PVCs. If it is set to <b>default</b> , the storage class name that is defined in <b>Default Storage Class</b> is used.	default
...	Leave everything up to but not including the <b>Product Name</b> parameters as default values.	default
Under CEM Configuration, configure <b>Product Name</b>	<ul style="list-style-type: none"> <li>For an eventing only installation, the <b>Product Name</b> is Event Management for IBM Multicloud Manager.</li> <li>For a full monitoring installation, the <b>Product Name</b> is IBM Cloud App Management for Multicloud Manger</li> </ul>	<b>IBM Cloud App Management for Multicloud Manger</b>
...	Leave everything up to but not including the <b>Cluster administrator Username</b> parameter as default values.	Default value
<b>Cluster administrator Username</b>	The username of your cluster administrator.	admin
...	Leave all remaining parameters as default values.	Default value

23. Click **Install** to deploy the `ibm-cloud-appmgmt-prod` Helm chart.

### Results

The Cloud App Management server is successfully installed with IBM Cloud Pak for Multicloud Management in full monitoring mode. If you want to optionally verify your installation, you can run the `collectContainerLogs.sh` script, which collects the installation logs and outputs them to a

diagnostic file. For more information about running this script, see [“Collecting the server logs for IBM Support” on page 1417](#).

## What to do next

Complete the following postinstallation steps to complete the OIDC registration:

1. In the console, click the menu in the upper left corner, and select **Monitor health > Helm Releases**, then select the release name that you gave for IBM Cloud App Management.
2. Scroll to the notes section at the bottom of your release and complete **step 3 (OIDC registration)** of the Notes section. Run the two **kubect1** commands that are displayed in step 3 in the Notes section.

**Note:** The **kubect1** commands are different in most cases, do not copy and paste them from this procedure. Instead, use the commands that are generated for you and your helm release. The first command registers with OIDC, and the second command adds two policy registrations. The following are example commands only.

```
COMMAND 1:
kubect1 exec -n kube-system -t `kubect1 get pods -l release=icam -n kube-system | grep "icam-ibm-cem-cem-users" | /
grep "Running" | head -n 1 | awk '{print $1}'` bash -- "/etc/oidc/oidc_reg.sh" "`echo $(
kubect1 get secret platform-oidc-credentials /
-o yaml -n kube-system | grep OAUTH2_CLIENT_REGISTRATION_SECRET: | awk '{print $2}'))`"
Registering IBM Cloud Event Management identity ...

Checking registration...
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
   Dload    Upload   Total   Spent    Left   Speed
100  135    0  135    0    0      537    0  --:--:--  --:--:--  --:--:--   544

* Client does not exist...
Registering client...

Done.

COMMAND 2:
kubect1 exec -n kube-system -t `kubect1 get pods -l release=icam -n kube-system | grep "icam-ibm-cem-cem-users" | /
grep "Running" | head -n 1 | awk '{print $1}'` bash -- "/etc/oidc/registerServicePolicy.sh" /
"`echo $(kubect1 get secret icam-cem-service-secret -o yaml -n kube-system | grep cem-
service-id: | awk '{print $2}'))" "`cloudctl tokens --access`"
Registering IBM Cloud Event Management identity ...
Checking service policy registration...
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
   Dload    Upload   Total   Spent    Left   Speed
100  15 100  15    0    0      8    0  0:00:01  0:00:01  --:--:--    8
Adding first policy registration
Adding second policy registration
Done
```

3. Now that the Cloud App Management server is successfully installed with IBM Cloud Pak for Multicloud Management, next you can deploy the ICAM klusterlet to monitor applications in your IBM Cloud Pak for Multicloud Management environment. For more information, see [“Installing the ICAM klusterlet on the managed cluster with Helm” on page 133](#) to deploy the ICAM klusterlet using its Helm chart or [“Installing the ICAM klusterlet on the managed cluster without helm” on page 135](#) for a non-Helm installation.

## Uninstalling IBM Cloud App Management

You can uninstall IBM Cloud App Management standalone or IBM Cloud App Management that was installed with IBM Cloud Pak for Multicloud Management using these steps.

### Before you begin

1. If you are uninstalling using the command line, you must install the Helm CLI. For instructions, see [Installing the Helm CLI \(helm\)](#).
2. You must delete the IBM Cloud Private service instance before you uninstall the Cloud App Management server. For more information, see [“Deleting the IBM Cloud Private service instance” on page 185](#).

## Procedure

Uninstall the Cloud App Management server from the console:

1. Log in to your console.
2. Click the hamburger menu in the upper left corner, and select **Helm Releases**.
3. From the list of helm releases, locate the Cloud App Management Helm release.  
If you know the name of your Cloud App Management Helm release, you can search for it by entering a keyword from the name.
4. Left-click on the **open and close list of options** menu (represented by three dots) for your Helm release and select **Delete**.

Uninstall the Cloud App Management server from the CLI:

5. Log in to your cluster. You only need to complete this step if you are uninstalling IBM Cloud App Management that was installed with IBM Cloud Pak for Multicloud Management. Ignore this step if you are uninstalling IBM Cloud App Management standalone.

```
cloudctl login -a my_cluster_URL -n kube-system --skip-ssl-validation
```

where

Where *my\_cluster\_URL* is the name that you defined for your cluster such as `https://cluster_address:443`. For future references to *masterIP*, use the value you are using for *cluster\_address*. A *cluster\_address* address example is: <https://icp-console.apps.organic-bullfrog-icp-mst.domain.com:443>.

6. Find the Helm chart that you want to uninstall from the list:

```
helm list --tls | grep ibm-cloud-appmgmt
```

7. Remove the Helm chart:

```
helm delete --purge --tls my_release_name
```

Where *my\_release\_name* is the name of your Cloud App Management Helm chart, such as `ibmcloudappmgmt`.

**Note:** Some CEM datalayer-cron jobs and pods might not be deleted. This is a known issue. Manually delete any remaining jobs or pods.

8. Delete the storage classes and persistent volume storage claims (PVCs) to release the claims on the persistent data store:

```
kubectl delete storageclass --selector release=my_release_name  
kubectl delete pvc --selector release=my_release_name --namespace my_namespace  
kubectl delete pv --selector release=my_release_name --namespace my_namespace
```

where *my\_namespace* is the namespace that the IBM Passport Advantage Archive (PPA) file is loaded to.

9. Delete secrets and the cluster image policy. You only need to run the second command: **kubectl delete clusterimagepolicy...** if you are uninstalling IBM Cloud App Management that was installed with IBM Cloud Pak for Multicloud Management. Ignore this command you are uninstalling IBM Cloud App Management standalone.

```
kubectl delete secrets --selector release=my_release_name --namespace my_namespace  
kubectl delete clusterimagepolicy --selector release=my_release_name --namespace my_namespace
```

10. Optional: Back up the data on the persistent storage directories that you created on the worker nodes.
11. Optional: You can safely remove the data from the persistent storage directories that you created on the worker nodes.



- Optional: You can remove the Cloud App Management image from IBM Cloud Private. For more information, see the [Removing an image from the console](#).

## Results

Cloud App Management server Helm chart is uninstalled. The storage configuration that was required for the installation is also deleted.

## Installing the ICAM klusterlet on the managed cluster with Helm

By installing the ICAM klusterlet on any managed cluster, you can use this ICAM klusterlet to configure and receive Kubernetes events from the cluster.

### Before you begin

The Cloud App Management server with IBM Cloud Pak for Multicloud Management must be installed on the hub cluster. For more information, select the specific installation scenario topic for your environment from “[Installing IBM Cloud App Management with IBM Cloud Pak for Multicloud Management](#)” on page 94.

You must import a managed cluster. For more information, see [Importing a target managed cluster to the hub cluster](#).

### About this task

Install the IBM Cloud CLI from the management console. Click **Menu > Command Line Tools > Cloud Private CLI** to download the installer by using a curl command. Copy and run the curl command for your operating system. For more information, see [Installing the IBM® Cloud Private CLI](#).

### Procedure

- Locate the ICAM klusterlet packages on [IBM Passport Advantage](#). Find the installation image by searching for it using its part number. Choose one of the following packages based on your platform:
  - For Linux® on Power®, choose the Multicloud Event Management Klusterlet for PlinuxLE (part number: CC4KXEN): `agent_ppa_2019.4.0_prod_ppc64le.tar.gz`
  - For Linux® x86\_64, choose Multicloud Event Management Klusterlet on AMD64 (part number: CC4KZEN): `agent_ppa_2019.4.0_prod_amd64.tar.gz`
- As an OpenShift administrator, log in to the OpenShift Container platform and log in to the Docker registry. Then, load the ICAM klusterlet PPA installation image file into the Docker registry:

```
cloudctl login -a my_cluster_URL -n my_namespace --skip-ssl-validation -u username -p password
docker login $(oc registry info) -u $(oc whoami) -p $(oc whoami -t)
cloudctl catalog load-archive --archive ppa_file --registry $(oc registry info)/my_namespace
```

Where `my_cluster_URL` is the name that you defined for your cluster such as `https://cluster_address:443`, `ppa_file` is the ICAM klusterlet PPA image file name, and `my_namespace` is `multicloud-endpoint`.

- In subsequent steps in the ICAM klusterlet configuration, you are prompted for the **MCM Fullname Override** value, which is the helm release name for the IBM Cloud Pak for Multicloud Management multicloud endpoint on the system. Run the following commands to determine this value. If you installed IBM Cloud Pak for Multicloud Management by using the Helm chart, the Helm release name is the first part of the pod name. In the following example, the Helm release name is `example`.

```
kubectl get pods --all-namespaces | grep klusterlet
kube-system example-ibm-mcm-klusterlet-klusterlet-66d565q5hnp 4/4 Running 0 2m
kube-system example-ibm-mcm-klusterlet-weave-scope-2f6m1 1/1 Running 0 2m
```

If you installed IBM Cloud Pak for Multicloud Management by using the cloudctl CLI instead of the Helm chart, run the following command on your IBM Cloud Pak for Multicloud Management multicluster endpoint to obtain the **MCM Fullname Override**.

```
kubectl get secrets -n multicluster-endpoint |grep hub-kubeconfig  
endpoint-connmgr-hub-kubeconfig Opaque 1 8d
```

The **MCM Fullname Override** value is endpoint-connmgr.

4. Log in to the management console of your target cluster.
5. Click **Catalog**.
6. Search for and select the **icam-clouddc-klusterlet** Helm chart.
7. Click **Configure**.
8. Configure the following configuration parameters:

**Helm release name**

Enter the helm release name for the ICAM klusterlet.

**Target namespace**

From the drop-down menu, select the namespace where IBM Cloud Pak for Multicloud Management is installed in [step 2](#).

**Target cluster**

The cluster on which icam-clouddc-klusterlet is being installed.

**Product License Acceptance**

Accept the product license by selecting the checkbox.

**Images Repository**

The image repository must be set to the registry name and namespace that the PPA was loaded to in [step 2](#).

**Note:** For OCP 4.2 or later environment that has openshift-image-registry route, use the internal docker image repository. Find the internal docker image repository by using the command:  
oc registry info --internal

**Images Prefix**

Leave this blank.

**Product Deployment Size**

From the drop-down menu, select the appropriate value.

**Replica Count**

Ignore this field.

**MCM Fullname Override**

The helm release name for the IBM Cloud Pak for Multicloud Management multicluster endpoint on the system. Use the value from [step 3](#).

**Alert target CRD**

Ignore this field.

9. Click **Install** to deploy the icam-clouddc-klusterlet Helm chart.

**What to do next**

1. Update the IBM Cloud Pak for Multicloud Management team for Cloud App Management. When a managed cluster is imported, the namespace corresponding to the managed cluster is created at the hub cluster. To activate monitoring of the managed cluster resources, add the namespace as a resource to a Cloud App Management team. On the console, click **Manage > Identity and Access > Teams**. For more information, see [Adding the Helm repository and namespace to a team](#).
2. On the console, click the **Incidents** menu to update the ICAM klusterlet with the correct webhook for Cloud App Management events.

3. Install and configure Cloud App Management agents and data collectors to collect data and metrics. For more information, see [“Deploying agents and data collectors for IBM Cloud Pak for Multicloud Management”](#) on page 140.

## Uninstalling the ICAM klusterlet that was installed with the Helm chart

Uninstall the ICAM klusterlet that you previously installed on a managed cluster using the Helm chart.

### Procedure

1. Purge the Helm release, which deletes the Kubernetes custom resource definition:

```
helm delete my_release_name --purge --tls
```

where *my\_release\_name* is the name of your IBM Cloud App Management Helm chart, such as `ibmcloudappgmt`.

2. Clean the Kubernetes customer resources and customer resource definition:

```
kubectl patch k8sdc k8sdc --namespace multicloudappgmt -p '{"metadata":{"finalizers":[]}}' \
--type=merge
kubectl delete crd k8sdc --namespace multicloudappgmt
```

3. Clean up the configuration secrets:

```
kubectl delete secret dc-secret --namespace multicloudappgmt
kubectl delete secret ibm-agent-https-secret --namespace multicloudappgmt
```

## Installing the ICAM klusterlet on the managed cluster without helm

After you install the Cloud App Management server, you can install the ICAM klusterlet to configure Prometheus and Kubernetes events. This procedure is for environments with no Helm installation, such as Red Hat OpenShift.

### Before you begin

The Cloud App Management server with IBM Cloud Pak for Multicloud Management must be installed on the hub cluster. For more information, select the specific installation scenario topic for your environment from [“Installing IBM Cloud App Management with IBM Cloud Pak for Multicloud Management”](#) on page 94.

You must import a managed cluster. For more information, see [Importing a target managed cluster to the hub cluster](#).

### Procedure

1. Locate the ICAM klusterlet packages on [IBM Passport Advantage](#). Find the installation image by searching for it using its part number. Choose one of the following packages depending on your platform:

- For Linux® on Power®, choose the Multicloud Event Management Klusterlet for PlinuxLE (part number: CC4KXEN): `agent_ppa_2019.4.0_prod_ppc64le.tar.gz`
- For Linux® x86\_64, choose Multicloud Event Management Klusterlet on AMD64 (part number: CC4KZEN): `agent_ppa_2019.4.0_prod_amd64.tar.gz`

2. Extract the Docker images:

```
tar xvf ppa_file images/
```

Where *ppa\_file* is the ICAM klusterlet PPA image file name from [“1”](#) on page 135.

3. Optional: Obtain your registry name, which you will use in the next step:

## Red Hat OpenShift 3.11:

```
# oc registry info
registry.default.svc:5000
```

For example:

```
docker load -i k8-monitor_APM_201912042117.tar.gz
docker tag k8-monitor:APM_201912042117 registry.default.svc:5000/multicluster-endpoint/k8-
monitor:APM_201912042117
docker push registry.default.svc:5000/multicluster-endpoint/k8-monitor:APM_201912042117
```

## Red Hat OpenShift 4.2

```
# oc get routes -n openshift-image-registry
NAME          HOST/PORT
PATH          SERVICES    PORT          TERMINATION  WILDCARD
default-route default-route-openshift-image-
registry.apps.icepick.os.fyre.ibm.com image-registry <all>      reencrypt   None
image-registry image-registry.openshift-image-
registry.svc image-registry image-registry 5000-tcp    reencrypt   None
```

For example:

```
docker load -i k8-monitor_APM_201912042117.tar.gz
docker tag k8-monitor:APM_201912042117 image-registry.openshift-image-registry.svc/
multicluster-endpoint/k8-monitor:APM_201912042117
docker push image-registry.openshift-image-registry.svc/multicluster-endpoint/k8-
monitor:APM_201912042117
```

### 4. Log in to your registry:

```
docker login $(oc registry info) -u $(oc whoami) -p $(oc whoami -t)
```

### 5. Load and push the following images to your Docker repository. The following images are for Linux® x86\_64. If your platform is Linux® on Power®, load and push these images to your Docker repository.

```
images/agentoperator_APM_201911270600.tar.gz
images/k8-monitor_APM_201912042117.tar.gz
images/k8sdc-operator_APM_201912042117.tar.gz
images/reloader_201910211632-multi-arch.tar.gz
```

a) Here is an example of how to load and push the k8-monitor Docker image to the repository:

```
docker load -i images/k8-monitor_APM_201912042117.tar.gz
docker tag k8-monitor:APM_201912042117 registry_name/multicluster-endpoint/k8-
monitor:APM_201912042117
docker push registry_name/multicluster-endpoint/k8-monitor:APM_201912042117
```

Where:

*registry\_name* is the name from step “3” on page 135

multicluster-endpoint is the target namespace on the cluster. The ICAM clusterlet is always installed in the same namespace where MCM clusterlet is installed, this is always multicluster-endpoint.

### 6. Create Docker **imagePullSecrets**:

```
kubectl create secret docker-registry my-registrykey --docker-server=$(oc registry info) --
docker-username=$(oc whoami) --docker-password=$(oc whoami -t) --docker-email=my_user_email
```

### 7. Deploy resources for the ICAM clusterlet:

a) Create the "custom resource definition":

```
kubectl apply -f k8monitor_crd.yaml
```

A template of the k8monitor\_crd.yaml file is shown here:

```
apiVersion: apiextensions.k8s.io/v1beta1
kind: CustomResourceDefinition
```

```

metadata:
  name: k8sdc.ibmcloudappmgmt.com
spec:
  group: ibmcloudappmgmt.com
  names:
    kind: K8sDC
    listKind: K8sDCList
    plural: k8sdc
    singular: k8sdc
  scope: Namespaced
  subresources:
    status: {}
  version: v1alpha1
  versions:
  - name: v1alpha1
    served: true
    storage: true

```

- b) Create the service account for the ICAM clusterlet:

```
kubectl apply -f service_account.yaml
```

A template of the `service_account.yaml` file is shown here:

```

apiVersion: v1
kind: ServiceAccount
metadata:
  name: icamklust
  namespace: multicluster-endpoint

```

- c) Apply the **imagePullSecrets** that you created in [step 5](#) to create the "icamklust" service account:

```
kubectl patch serviceaccount icamklust -p '{"imagePullSecrets": [{"name": "<my-registrykey>"}]}' -n multicluster-endpoint
```

- d) Bind cluster-admin with the "icamklust" service account. For OpenShift monitoring, you must create a **ClusterRoleBinding**.

```

oc create clusterrolebinding my_cluster_role_binding_name \
  --clusterrole=cluster-admin \
  --serviceaccount=multicluster-endpoint:icamklust -n multicluster-endpoint

```

where *my\_cluster\_role\_binding\_name* is a new cluster role binding name

multicluster-endpoint is the target namespace on the cluster (the project name in OpenShift). The ICAM clusterlet is always installed in the same namespace where the MCM clusterlet is installed, this is always multicluster-endpoint.

- e) Bind cluster-admin with the "default":

```

oc create clusterrolebinding my_cluster_role_binding_name_default --clusterrole=cluster-admin \
  --serviceaccount=multicluster-endpoint:icamklust -n multicluster-endpoint

```

where *my\_cluster\_role\_binding\_name\_default* is a new cluster role binding name.

- f) Create the k8sdc-operator deployment:

```
kubectl apply -f k8sdc-operator.yaml
```

A template of the `k8sdc-operator.yaml` file is shown here:

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: k8sdc-operator
  namespace: multicluster-endpoint
spec:
  replicas: 1
  selector:
    matchLabels:
      name: k8sdc-operator
  template:

```

```

metadata:
  labels:
    name: k8sdc-operator
spec:
  serviceAccountName: icamklust
  containers:
  - name: k8sdc-operator
    # Replace this with the built image name
    image: registry_name/multicluster-endpoint/k8sdc-operator:APM_202002202250
    imagePullPolicy: Always
    env:
      - name: WATCH_NAMESPACE
        valueFrom:
          fieldRef:
            fieldPath: metadata.namespace
      - name: POD_NAME
        valueFrom:
          fieldRef:
            fieldPath: metadata.name
      - name: OPERATOR_NAME
        value: "k8sdc-operator"

```

g) Create the reloader deployment:

```
kubectl apply -f icam-reloader.yaml
```

A template of the `icam-reloader.yaml` file is shown here:

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: icam-reloader
  namespace: multicluster-endpoint
spec:
  replicas: 1
  revisionHistoryLimit: 2
  selector:
    matchLabels:
      name: icam-reloader
  template:
    metadata:
      labels:
        name: icam-reloader
    spec:
      securityContext:
        runAsNonRoot: true
        runAsUser: 1000
      containers:
      - env:
        - name: KUBERNETES_NAMESPACE
          valueFrom:
            fieldRef:
              apiVersion: v1
              fieldPath: metadata.namespace
          # Replace this with the built image name
          image: registry_name/multicluster-endpoint/reloader:201910211632-multi-arch
        resources:
          limits:
            cpu: "500m"
            memory: "100Mi"
          requests:
            cpu: "50m"
            memory: "50Mi"
          imagePullPolicy: Always
          name: icam-reloader

```

where `registry_name` is the name from step “3” on page 135

h) Create the agentoperator deployment:

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: agentoperator
  namespace: multicluster-endpoint
spec:
  replicas: 1
  selector:
    matchLabels:

```

```

    name: agentoperator
  template:
    metadata:
      labels:
        name: agentoperator
    spec:
      serviceAccountName: icamklust
      containers:
        - name: agentoperator
          # Replace this with the built image name
          image: registry_name/multicluster-endpoint/agentoperator:APM_202002241310
          imagePullPolicy: Always
          command:
            - agentoperator
          args:
            - "--image-repo=registry_name/multicluster-endpoint"
            - "--image-prefix="
          env:
            - name: POD_NAME
              valueFrom:
                fieldRef:
                  fieldPath: metadata.name
            - name: OPERATOR_NAME
              value: "agentoperator"
          volumeMounts:
            - name: klusterlet-config
              mountPath: /opt/klusterlet
      volumes:
        - name: klusterlet-config
          secret:
            secretName: endpoint-connmgr-hub-kubeconfig

```

```
kubectl apply -f agentoperator.yaml
```

Validate the deployments that you created in the previous steps:

8. After the installation script is finished, run the following commands:

```

kubectl get deployment k8sdc-operator --namespace=multicluster-endpoint
kubectl get deployment icam-reloader --namespace=multicluster-endpoint
kubectl get deployment agentoperator --namespace=multicluster-endpoint

```

## What to do next

1. Update the IBM Cloud Pak for Multicloud Management team for Cloud App Management. When a managed cluster is imported, the namespace corresponding to the managed cluster is created at the hub cluster. To activate monitoring of the managed cluster resources, add the namespace as a resource to a Cloud App Management team. On the console, click **Manage > Identity and Access > Teams**. For more information, see [Adding the Helm repository and namespace to a team](#).
2. On the console, click the **Incidents** menu to update the ICAM klusterlet with the correct webhook for Cloud App Management events.
3. Install and configure Cloud App Management agents and data collectors to collect data and metrics. For more information, see [“Deploying agents and data collectors for IBM Cloud Pak for Multicloud Management” on page 140](#).

## Uninstalling the ICAM klusterlet that was installed without the Helm chart

Uninstall the ICAM klusterlet that you previously installed on a managed cluster without the Helm chart.

### Procedure

1. Delete the agentoperator:

```
oc delete deployment agentoperator -n multicluster-endpoint
```

2. Delete the Kubernetes custom resource instance:

```
oc delete k8sdc k8sdc-cr -n multicluster-endpoint
```

3. Delete the Kubernetes operator:

```
oc delete deployment k8sdc-operator -n multicluster-endpoint
```

4. Delete the reloader:

```
oc delete deployment icam-reloader -n multicluster-endpoint
```

5. Find the Kubernetes custom resource definition:

```
oc get crd |grep k8sdc
```

6. Delete the Kubernetes custom resource definition.

```
oc delete crd kubernetes_crd
```

where *kubernetes\_crd* is the Kubernetes custom resource definition that you retrieved in [step 5](#).

7. Delete the clusterrolebinding for service account:

```
oc delete clusterrolebinding my_cluster_role_binding_name
```

8. Delete the service account:

```
oc delete sa my_service_account_name -n multicluster-endpoint
```

where *my\_service\_account\_name* is the service account name that you specified when you installed the ICAM klusterlet without the Helm chart.

when setting up the service account during helm-free installation of the klusterlet

9. Delete the Docker imagePullSecrets:

```
oc delete secret my_docker_image_secret -n multicluster-endpoint
```

10. Delete the dc-secret:

```
oc delete secret dc-secret -n multicluster-endpoint
```

11. Delete the ibm-agent-https-secret secret:

```
oc delete secret ibm-agent-https-secret -n multicluster-endpoint
```

## Deploying agents and data collectors for IBM Cloud Pak for Multicloud Management

After you deploy the ICAM klusterlet, you can install and configure agents and data collectors.

### Before you begin

The ICAM klusterlet must be deployed. For more information, see [“Installing the ICAM klusterlet on the managed cluster with Helm”](#) on page 133.

### About this task

Kubernetes data collector is auto-configured when you deploy the ICAM klusterlet. The other agents and data collectors still need manual configuration.

You must download the data collector configuration packs to configure agents and data collectors to communicate with the Cloud App Management server.

### Procedure

To download the configuration packs in IBM Multicloud Manager, complete the following steps:

1. Login to the console.
2. Select **Event Management** from the menu on the upper left of the window.
3. Click **Integrations** on the IBM Cloud App Management **Administration** page.
4. Click **New Integration**.



5. Click **Configure** under **ICAM Data collectors** or **ICAM Agents** depending on which one you want to configure.
6. Download the configuration package.

### What to do next

For further instructions about deploying ICAM agents and data collectors, see the following topics:

- [Chapter 13, “Deploying ICAM Agents,” on page 193.](#)
- [Chapter 15, “Deploying ICAM Data Collectors,” on page 555.](#)
- IBM Cloud App Management provides the Unified Agent, which is a framework of plug-ins to collect, process, aggregate, and write metrics to your Cloud App Management environment. It is based on Telegraf. By deploying the Unified Agent, you can receive OpenTracing workloads such as Jaeger and Zipkin, monitor NGINX and Redis workloads, IBM API Connect®, IBM App Connect Enterprise, and IBM MQ. To learn how to deploy the Unified Agent, see [Chapter 16, “Deploying Unified Agent,” on page 645.](#)

## Offline: Installing IBM Cloud App Management stand-alone on Red Hat OpenShift

These steps explain how to install the IBM Cloud App Management server on Red Hat OpenShift.

### Before you begin

1. Download and install Red Hat OpenShift V3.11.
2. Download and install the IBM Cloud Private V3.2.1 for Red Hat Enterprise Linux OpenShift (64-bit) Docker package (Part number: CC3KREN) from [IBM Passport Advantage®](#). Find the installation image by searching for it using its part number. Then, install its most recent fix pack from [IBM Fix Central](#). Search for IBM Cloud Private to find its associated fix packs.
3. Install the Helm CLI. For instructions, see [Installing the Helm CLI \(helm\)](#).
4. Install the cloudctl command-line CLI. For instructions, see [Installing cloudctl](#).
5. It is best practice to install the Cloud App Management server in a new, non-default namespace. There was a new `limitrange` resource added in IBM Cloud Private V3.2.1 for the default namespace. Installing the Cloud App Management server into the default namespace of IBM Cloud Private V3.2.1 can be impacted by this limit range. If you are not using a shared cluster, you can delete the `limitrange` resource in the default namespace by running the following command:

```
kubectl delete limitrange default-limit -n default
```

6. You must onboard LDAP users. For more information, see [“Onboarding LDAP users” on page 94.](#)
7. Review the optimization for performance topics. For more information, see [“Optimizing performance” on page 156.](#)
8. The Cloud App Management server uses the UIDs 100, 1000 and 1001 and GIDs 100, 1000 and 1001. To avoid any issues with file ownership or permissions, you can create users and groups for each UID and GID. For example, create a username "icam" with UID 100 and a group "icamgrp" with GID 100. Create the users and groups for UID 1000 and 1001 and GID 1000 and 1001. You can choose any user and group names when you create them. While not required, it is considered best practice to create these users and groups before you install to help avoid any confusion with file and process ownership. If any users or groups exist with the UIDs or GIDs 100, 1000, or 1001 then you might observe files and processes that are owned by those users and groups.
9. Configure the Elasticsearch `vm.max_map_count` parameter on all worker nodes by completing the following steps:
  - a. For Elasticsearch, you must set a kernel parameter on all worker nodes. These nodes are identified when you configure the persistent storage later in this procedure. You must set `vm.max_map_count` to a minimum value of 1048575. Set the parameter by using the `sysctl`

command to ensure that the change takes effect immediately. Run the following command on each worker node:

```
sysctl -w vm.max_map_count=1048575
```

- b. You must also set the **vm.max\_map\_count** parameter in the `/etc/sysctl.conf` file to ensure that the change is still in effect after the node is restarted.

```
vm.max_map_count=1048575
```

## Procedure

1. Locate and download the Cloud App Management server installation image file `icam_ppa_2019.4.0_prod.tar.gz` (Part number: CC4KNEN) from [IBM Passport Advantage](#). Find the installation image by searching for it using its part number. For more information on the IBM Cloud App Management components and their part numbers, see [“Part numbers” on page 71](#).
2. As the administrator, log in to the management console. Run the following command:

```
cloudctl login -a my_cluster_URL -n my_namespace --skip-ssl-validation
```

Where `my_cluster_URL` is the name that you defined for your cluster such as `https://cluster_address:443` and `my_namespace` is the namespace where you are installing your Cloud App Management server. For future references to `masterIP`, use the value you are using for `cluster_address`. A `cluster_address` address example is: `icp-console.apps.organic-test-icp-mst.domain.com`.

3. As an OpenShift administrator, log in to the OpenShift Container Platform:

```
oc login
```

4. Log in to the Docker registry:

```
docker login $(oc registry info) -u $(oc whoami) -p $(oc whoami -t)
```

**Note:** For OpenShift V4.2 or later environment, if you get a `x509: certificate signed by unknown authority` error, you must complete the step 6 to 10 also. Ignore these steps for OpenShift V3.11.

5. Create an installation directory and download the IBM Passport Advantage file: `icam_ppa_2019.4.0_prod.tar.gz` to the installation directory.

```
mkdir -p install_dir
```

6. Move the IBM Passport Advantage Cloud App Management server installation image file: `icam_ppa_2019.4.0_prod.tar.gz` to the `install_dir` directory and `cd` to the installation directory:

```
mv app_mgmt_server_2019.4.0.tar.gz install_dir/  
cd install_dir
```

**Note:** The Cloud App Management server installation image file is approximately 14 GB. Ensure that you have enough free space to store the file.

7. Extract the Helm chart from the `icam_ppa_2019.4.0_prod.tar.gz` installation image file and decompress it into the installation directory:

```
tar -xvf ./icam_ppa_2019.4.0_prod.tar.gz charts  
tar -xvf ./charts/ibm-cloud-appmgmt-prod-1.6.0.tgz
```

8. Load the image archive into the catalog. This process can take 10 - 30 minutes.

```
cloudctl catalog load-archive --archive ./icam_ppa_2019.4.0_prod.tar.gz --registry $(oc registry info)/my_namespace
```

where *my\_namespace* is the namespace you logged in to `cloudctl` with earlier and it is where you are installing your Cloud App Management server.

9. Create local directories on the selected IBM common services worker nodes where you want the statefulsets to run. Record the IP address of each worker node where the directories are created. You need these IP addresses when you run the `prepare-pv.sh` script in the next step. For more information about optimization performance, see [“Optimizing performance” on page 156](#).

**Note:** The Cassandra persistent volume should reside on its own disk and dedicated worker node. Avoid creating the other persistent volumes on the same worker node and the same persistent volume with Cassandra.

**Note:** The ZooKeeper and Kafka persistent volumes should reside on the same worker node.

**Note:** The CouchDB and Datalayer persistent volumes should reside on the same worker node.

- Cassandra persistent storage:

```
ssh root@cassandraNode
mkdir -p /k8s/data/cassandra
```

- ZooKeeper persistent storage:

```
ssh root@zookeeperNode
mkdir -p /k8s/data/zookeeper
```

- Kafka persistent storage:

```
ssh root@kafkaNode
mkdir -p /k8s/data/kafka
```

- CouchDB persistent storage:

```
ssh root@couchdbNode
mkdir -p /k8s/data/couchdb
```

- Datalayer persistent storage:

```
ssh root@datalayerNode
mkdir -p /k8s/data/datalayer
```

- Elasticsearch persistent storage

```
mkdir -p /k8s/data/elasticsearch
```

10. Create the storage classes and persistent volume yam files for each of the following Cloud App Management statefulsets by running the `ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh` script. The use of persistent volumes that are backed by local disks or local partitions is highly recommended.

**Note:** For more information about storage, after you select the target **ibm-cloud-appmgmt-prod** chart from the catalog, select the **Overview** tab if it is not already selected. Next, scroll down the Storage section on the IBM Cloud App Management page on the right. For more information about how to configure persistent storage, see the following topics: [“Planning hardware and sizing” on page 77](#), [Understanding Kubernetes storage](#), [Planning a storage solution](#), and [Planning persistent storage](#).

```
Usage: ./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh
--releaseName <name>           Release name (default is ibmcloudappmgmt)
--size0_amd64                   Install as size0 on amd64 (minimum resource
requirements) - if omitted, specify each size in parameters
--size0_ppc64le                 Install as size0 on ppc64le (minimum resource
requirements) - if omitted, specify each size in parameters
--size1_amd64                   Install as size1 on amd64 (standard resource
requirements) - if omitted, specify each size in parameters
--size1_ppc64le                 Install as size1 on ppc64le (standard resource
requirements) - if omitted, specify each size in parameters
*Required flags for local storage:
```

```

--local                                Use local persistent volume storage
--CassandraNodes <worker_node_IP>    Worker node(s) for Cassandra. For multiple
Cassandra use a quoted, space separated list.
--KafkaNodes <worker_node_IP>        Worker node(s) for Kafka. For multiple Kafka use
a quoted, space separated list.
--ZookeeperNodes <worker_node_IP>    Worker node(s) for Zookeeper. For multiple
Zookeeper use a quoted, space separated list.
--CouchDBNodes <worker_node_IP>      Worker node(s) for CouchDB. For multiple CouchDB
use a quoted, space separated list.
--DatalayerNodes <worker_node_IP>     Worker node(s) for Datalayer. For multiple
Datalayer use a quoted, space separated list.
--ElasticsearchNodes <worker_node_IP> Worker node(s) for Elasticsearch. For multiple
Elasticsearch use a quoted, space separated list.

*Optional storage directory paths for local storage:
--CassandraDir <directory>           Local system directory for Cassandra (default
is /k8s/data/cassandra)
--CassandraBackupDir <directory>     Local system directory for Cassandra backups
(default is /k8s/data/cassandra_backup)
--KafkaDir <directory>               Local system directory for Kafka (default is /k8s/
data/kafka)
--ZookeeperDir <directory>          Local system directory for Zookeeper (default
is /k8s/data/zookeeper)
--CouchDBDir <directory>            Local system directory for CouchDB (default is /k8s/
data/couchdb)
--DatalayerDir <directory>          Local system directory for Datalayer (default
is /k8s/data/datalayer)
--ElasticsearchDir <directory>      Local system directory for Elasticsearch (default
is /k8s/data/elasticsearch)

*Optional storage class name flags for local storage:
--CassandraClass <className>         Storage class name for Cassandra (default is
<release_name>-local-storage-cassandra)
--CassandraBackupClass <className>    Storage class name for Cassandra backups
(default is <release_name>-local-storage-cassandra-backup)
--KafkaClass <className>             Storage class name for Kafka (default is
<release_name>-local-storage-kafka)
--ZookeeperClass <className>         Storage class name for Zookeeper (default is
<release_name>-local-storage-zookeeper)
--CouchDBClass <className>           Storage class name for CouchDB (default is
<release_name>-local-storage-couchdb)
--DatalayerClass <className>         Storage class name for Datalayer (default is
<release_name>-local-storage-datalayer)
--ElasticsearchClass <className>     Storage class name for Elasticsearch (default
is <release_name>-local-storage-elasticsearch)

*Required flags for vSphere storage:
--vSphere                             Use vSphere provisioned storage (requires existing
vSphere storage class)

*Optional storage size flags for local and vSphere storage:
--CassandraSize <size>               Size of persistent volume for Cassandra
(default size0_amd64 and size0_ppc64le is 50Gi)
--CassandraBackupSize <size>         Size of persistent volume for Cassandra
backups (default size0_amd64 and size0_ppc64le is 50Gi)
--KafkaSize <size>                   Size of persistent volume for Kafka (default
size0_amd64 and size0_ppc64le is 5Gi)
--ZookeeperSize <size>               Size of persistent volume for Zookeeper
(default size0_amd64 and size0_ppc64le is 1Gi)
--CouchDBSize <size>                 Size of persistent volume for CouchDB
(default size0_amd64 and size0_ppc64le is 5Gi)
--DatalayerSize <size>               Size of persistent volume for Datalayer
(default size0_amd64 and size0_ppc64le is 5Gi)
--ElasticsearchSize <size>           Size of persistent volume for Elasticsearch
(default size0_amd64 and size0_ppc64le is 5Gi)

```

**Note:** The --local and --vSphere are incompatible. You must choose one or the other. Local storage is highly recommended. If you don't choose either the --local or --vSphere options, the prepare-pv.sh defaults to local storage.

```
kubect1 get nodes
```

You don't need to use all the parameters, here are some example scenarios with the typical parameters used:

Scenario 1: No High Availability with a total of 2 VMs. 1 for Cassandra, 1 for the remaining statefulsets:

```
ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh --size1_amd64
--cassandraNode "worker01" --zookeeperNode "worker02"
--kafkaNode "worker02" --couchdbNode "worker02" --datalayerNode "worker02" --
elasticsearchNode "worker02"
```

Scenario 2: High Availability with a total of 6 VMs. 3 for Cassandra, 3 for the remaining statefulsets:

```
ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh --size1_amd64
--cassandraNode "worker01 worker02 worker03" --zookeeperNode "worker04 worker05 worker06"
--kafkaNode "worker04 worker05 worker06" --couchdbNode "worker04 worker05 worker06"
--datalayerNode "worker04 worker05 worker06" --elasticsearchNode "worker04 worker05
worker06"
```

Scenario 3: High Availability with a total of 9 VMs. 6 for Cassandra, 3 for the remaining statefulsets:

```
ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh
--size1_amd64 --cassandraNode "worker01 worker02 worker03 worker07 worker08 worker09"
--zookeeperNode "worker04 worker05 worker06" --kafkaNode "worker04 worker05 worker06" --
couchdbNode "worker04 worker05 worker06"
--datalayerNode "worker04 worker05 worker06" --elasticsearchNode "worker04 worker05
worker06"
```

11. The yaml files that make up your storage classes and persistent volumes are now created. Next, you must create the Kubernetes resources that will use them. Navigate to the `ibm-cloud-appmgmt-prod/ibm_cloud_pak/yaml/` directory, and create the storage classes and persistent volumes in there:

```
cd ibm-cloud-appmgmt-prod/ibm_cloud_pak/yaml/
oc create -f . -n my_namespace
cd ../../.. # this should move you back to the install_dir directory
```

12. Run the `ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/pre-install.sh` script.

The `pre-install.sh` completes many functions. At a minimum, the script creates a `values.yaml` file that you can use to override the default values set in the Cloud App Management Helm chart. It can also be used to enable TLS encryption between the agents and the Cloud App Management server. Refer to the following code snippet to check which flags are required and optional. Examples are given after the code snippet.

```
Usage ./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/pre-install.sh
Use this script to perform preparation tasks that require admin permissions before IBM
Cloud AppMgmt is installed.

  *Required flags
  --accept                Accepts license agreement(s)
  --https                 Install with HTTPS enabled (HTTPS is always
enabled in Advanced offering)
  --advanced              Install as ADVANCED offering (omit this
parameter will install as Base offering)
  --releaseName <name>   Release name (default is ibmcloudappmgmt)
  --masterAddress <FQDN> Fully qualified domain name (FQDN) for the ICP
Master.
                        In a highly available environment, this would be
the FQDN of the HAProxy or load balancer for ICP.
  --masterPort <int>     The port of the ICP master. On OpenShift the
default port is 443, which will need to be specified here. (default is 8443)
  --ingressPort <int>   The ingress port used to access the ICP console.
(default is 443)
  --proxyIP <IP>        IP address for ICP Proxy.
                        In a highly available environment, this would be
the IP address of the HAProxy or load balancer for ICP.
  --proxyFQDN <FQDN>    Fully qualified domain name (FQDN) for the ICP
Proxy.
                        In a highly available environment, this would be
the FQDN of the HAProxy or load balancer for ICP.
  --namespace <name>    Namespace (default is default)
  --clusterCAdomain <name> ICP cluster domain name, default is mycluster.icp
                        This value, combined with the values provided
for --repositoryPort and --namespace, will determine the imageRepository where ICAM will
```

look for images, e.g. with --clusterCADomain mycluster.icp --repositoryPort 8500 and --namespace icam, you will end up with an imageRepository of mycluster.icp:8500/icam  
 --repositoryPort <int> The port of the image repository. On OpenShift the default port is 5000, which will need to be specified here. (default is 8500)

\*Optional - Email setup:

--emailtype <smtp api>	Type of email, either smtp or api
--emailfrom <emailAddress>	Email address to show on sent mail as from
--smtp host <hostname>	SMTP hostname
--smtpport <port>	SMTP port
--smtpuser <user>	SMTP user
--smtp pass <password>	SMTP password
--smtpauth <true false>	User authentication required for SMTP connection (default is true)
--smtprejectunauthorized <true false>	Set this to false to allow self signed certificates when connecting via TLS, true enforces TLS authorization checking (default is true)
--apikey <key>	API key file

\*Optional - High availability and horizontal scale settings

--minReplicasHPAs <int>	The minimum number of replicas for each deployment, controlled by HPAs
--maxReplicasHPAs <int>	The maximum number of replicas for each deployment, controlled by HPAs
--kafkaClusterSize <int>	The number of Kafka replicas (the replication factor for Kafka topics will be set to this value, up to a max of 3)
--zookeeperClusterSize <int>	The number of Zookeeper replicas (all Zookeeper data is replicated to all zookeeper nodes)
--couchdbClusterSize <int>	The number of CouchDB replicas (the CouchDB data replication defaults to 3, even if the cluster has 1 or 2 nodes)
--datalayerClusterSize <int>	The number of Datalayer replicas (the datalayer relies on Kafka and internal jobs for handling data replication)
--elasticsearchClusterSize <int>	The number of Elasticsearch replicas (the number of replica shards is determined from the number of Elasticsearch instances)
--redisServerReplicas <int>	The number of redis server replicas. Defaults to 1
--cassandraClusterSize <int>	The number of Cassandra replicas (the replication factor for Cassandra keyspaces will be set to this value, up to a max of 3)
--cassandraUsername <string>	The username Cassandra will use. You must use a username other than 'cassandra'. If left unset, the default cassandra credentials will be used.

\*Optional - Other

--metricSummarization <string>	Enables or disables metric summarization. Set to 'true' or 'false'. Defaults to 'false' if not specified.
--metricC8Rep <replication_string>	The replication string for the metric data (default is "{class:'SimpleStrategy','replication_factor':X}", where X is the cassandraClusterSize up to 2)
--openttC8Rep <int>	The replication factor for the Open Transaction Tracking data (default is to match cassandraClusterSize up to 2)
--metricKafkaRep <int>	The replication factor for the metric Kafka data (default is to match kafkaClusterSize up to 2)
--useTLSCertsJob <true false>	Enables or disables creating Ingress TLS Certificates using Kubernetes Job. Set to 'true' or 'false'. Defaults to 'false' if not specified.

Example of install an Advanced offering with HTTPS enabled on Openshift:

```
./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/pre-install.sh --accept \
--releasename icam \
--namespace default \
--masteraddress x.xx.xx.xx \
--proxyip x.xx.xx.xx \
--proxyfqdn proxy.example.com \
--clustercadomain docker-registry.default.svc \
--repositoryPort 5000 \
--advanced \
--cassandraUsername customCassandraSuperuser \
--masterPort 443 \
--ingressPort 443
```

Example of install an Advanced offering with HTTPS enabled on Openshift, using high availability:

```
./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/pre-install.sh --accept \
--releasename icam \
--namespace default \
--masteraddress haproxy.example.com \
--proxyip x.xx.xx.xx \
--proxyfqdn haproxy.example.com \
--clustercadomain docker-registry.default.svc \
--repositoryPort 5000 \
--advanced \
--minreplicashpas 2 \
```

```
--maxreplicashpas 3 \
--kafkaclustersize 3 \
--zookeeperclustersize 3 \
--couchdbclustersize 3 \
--datalayerclustersize 3 \
--cassandraclustersize 3 \
--redisServerReplicas 3 \
--cassandraUsername customCassandraSuperuser \
--masterPort 443 \
--ingressPort 443
```

13. Optional: If you want to change the raw metric retention period from the default 8 days, use the `--set` option when you are running the **helm install** command in step “17” on page 147.

```
--set global.metric.retention.rawMaxDays=2
```

where 2 represents the number of days to retain and can be a whole number 2 - 32. Any value beyond 32 days is not recommended and can compromise Cloud App Management performance.

The following example shows a Helm installation command that sets the data retention to 15 days:

```
helm install --name ibmcloudappgmt --values ibmcloudappgmt.values
--set global.metric.retention.rawMaxDays=15 my_install_dir/ibm-cloud-appgmt-
prod-1.6.0.tgz --tls
```

14. Optional: If you want to enable storage of summarized metric data, specify the following option on the helm install command:

```
--set global.metric.summary.enabled=true
```

For more information, see “Data retention and summarization” on page 765.

15. Change to the `clusterAdministration` directory and run the `createSecurityClusterPrereqs.sh` script to create the `SecurityContextConstraint` for the Cloud App Management resources to use:

```
cd ~/install_dir/ibm-cloud-appgmt-prod/ibm_cloud_pak/pak_extensions/pre-install/
clusterAdministration
./createSecurityClusterPrereqs.sh
```

16. Change to the `install_dir` directory.

```
cd ~/install_dir
```

17. Deploy the Cloud App Management server Helm chart using the Helm CLI.

```
helm install --name my_release_name --namespace my_namespace --values
my_release_name.values.yaml \
charts/ibm-cloud-appgmt-prod-1.6.0.tgz --tls
```

18. After the Helm chart is successfully deployed, run the `ibm-cloud-appgmt-prod/ibm_cloud_pak/pak_extensions/post-install-setup.sh` script to complete administrative tasks necessary to access the IBM Cloud App Management dashboard. Run the script without any parameters first to see which parameters are required and optional. For example:

```
Usage: ./ibm-cloud-appgmt-prod/ibm_cloud_pak/pak_extensions/post-install-setup.sh

--releaseName <name>          Release name, default of ${default_release}"
--namespace <name>           Namespace, default of ${namespace}"
--instanceName <name>        Name for the serviceinstance, default of ${instance_name}"

[ --advanced ]                Choose Advanced offering ( omit this parameter will chose
Base offering )"
[ --noLog ]                   Do not log to ${log_file}"
[ --tenantID <UUID> ]        The TenantID of the new serviceinstance, default is random"

"example: for Base offering"
./ibm-cloud-appgmt-prod/ibm_cloud_pak/pak_extensions/post-install-setup.sh \
--releaseName ${default_release} --instanceName ${instance_name} --namespace ${namespace}"

"example: for Advanced offering"
./ibm-cloud-appgmt-prod/ibm_cloud_pak/pak_extensions/post-install-setup.sh \
--releaseName ${default_release} --instanceName ${instance_name} --namespace ${namespace} --
```

advanced"

**Note:** After you run the `post-install-setup.sh` script, sometimes a `Connection timed out` error can be displayed. For more information about solving this issue, see [“The post-install-setup.sh script runs with the timeout error on Red Hat OpenShift environment”](#) on page 1406.

## Results

The Cloud App Management server is successfully installed. If you want to optionally verify your installation, you can run the `collectContainerLogs.sh` script, which collects the installation logs and outputs them to a diagnostic file. For more information about running this script, see [“Collecting the server logs for IBM Support”](#) on page 1417.

## What to do next

1. Start the service instance and access the Cloud App Management console. For more information, see [“Starting the Cloud App Management UI”](#) on page 176.
2. Deploy ICAM agents. For more information, see [Chapter 13, “Deploying ICAM Agents,”](#) on page 193.

# Offline: Installing IBM Cloud App Management stand-alone on IBM Cloud Private

These steps explain how to install the Cloud App Management server in an IBM Cloud Private environment.

## Before you begin

To ensure that your server deployment is successful, you must first complete the required planning tasks on your system. Determine the hardware requirements and storage type for your environment. For more information, see [Chapter 7, “Planning your deployment,”](#) on page 69.

## About this task

Deploying the server includes the following main steps:

1. **Optional:** Review the optimization performance topics. For more information, see [“Optimizing performance”](#) on page 156
2. **Optional:** If you want to deploy Cloud App Management as a highly available set of services, providing redundancy, before you begin the installation, review [“Planning for a high availability installation”](#) on page 148.
3. **Optional:** Configure a custom certificate. For more information, see [“Configuring a custom server certificate”](#) on page 153.
4. Install the Cloud App Management server. For instructions, see [“Installing IBM Cloud App Management stand-alone on IBM Cloud Private”](#) on page 163.
5. **Optional:** Move to a custom namespace. For more information, see [“Moving to a custom namespace”](#) on page 172.
6. **Optional:** Validate the Cloud App Management server deployment is successful. For more information, see [“Validating the Cloud App Management server deployment”](#) on page 173.
7. Complete the postinstallation task of creating your service instance. For more information, see [“Creating your service instance”](#) on page 174.
8. Access the Cloud App Management console. For more information, see [“Starting the Cloud App Management UI”](#) on page 176.

## Planning for a high availability installation

To deploy Cloud App Management server as a highly available microservice, multiple instances (replicas) of the stateful and stateless services are required to provide redundancy.



## Stateless services

Horizontal Pod Autoscaler (HPA) provides scalability for stateless services as described here: [“Scaling stateless and stateful services”](#) on page 83. The HPAs have a minimum and maximum number of replicas they can scale to. This value is determined by the **minReplicasHPAs** and **minReplicasHPAs** global parameters, which are specified when you run `pre-install.sh` script. By adjusting the **minReplicasHPAs** to a number greater than 1, you can provide a level of high availability by ensuring that multiple instances are deployed by Kubernetes.

After installation, you can adjust these values globally by running the `pre-install.sh` script and running a helm upgrade. Alternatively, you can edit specific services by editing the HPA either through the CLI or the IBM Cloud Private UI. To see a list of stateless services, see [“Scaling stateless and stateful services”](#) on page 83.

## Stateful services

There are five stateful services that are needed for the operations of Cloud App Management server: Cassandra, Kafka, ZooKeeper, CouchDB, and Datalayer. If a disk or system failure occurs on one of the stateful services instances, the other remaining replicas provide resiliency and keep the services operating. The workloads are automatically spread across the stateful service instances, distributing the work and providing resiliency. Since data is spread evenly between replicas of the statefulset, each PersistentVolume (PV) needs to be the same size. The following section describes the resiliency planning considerations for the stateful services:

### Cassandra

Cassandra is a distributed database with no single point of failure. When data is stored in a multinode environment, a hash is used to decide which node gets which data, spreading the data evenly across the cluster. With 3 nodes and a replication factor of 3, each node has a copy of 100% of the data. With 3 nodes and a replication factor of 2, each node has 67% of the data. As you expand out from there, each node has a decreasing percentage of the data. Since each node is responsible for less data, adding more nodes increases your cluster capacity. When a node goes offline, the remaining nodes pick up the slack. For more information, view the [Cassandra](#) documentation.

The Cassandra cluster size is set via the `--cassandraClusterSize` flag in the `pre-install.sh` script. This controls the **global.cassandraNodeReplicas** yaml value. Replication is set on a per keyspace basis. When you run the `pre-install.sh` script, if you select 3 or more for `--cassandraClusterSize`, the keyspaces are configured as follows:

- The Topology (janusgraph keyspace) and Events (datalayer keyspace) keyspaces will be configured with a replication factor of 3.

- The metric data (metricdb keyspace) and Open Tracing (jaeger\_v1\_opentt keyspace) will be configured with a replication factor of 2.

The Metric data and Open Tracing keyspaces have larger levels of traffic and are less system critical than the Topology and Event spaces, making the tradeoff acceptable. To override this, set other values in the `--metricC8Rep` and `--openttC8Rep` parameters when you run the `pre-install.sh` script. This will set the **global.metricC8Rep** and **global.openttC8Rep** yaml values.

Because the Topology and Event replication factors is 3, for a high availability environment, 3 or more Cassandra instances are required. This is because with replication greater than 2, Cassandra needs a quorum of members to keep consistency. A Cassandra quorum is defined as half plus one. If the data replication factor is set to 3, a quorum would be 2. A quorum of 2 nodes would be  $1+1=2$ , which would mean you would not be able to lose any members. The metric and open tracing data is queried with consistency of "ONE", not "QUORUM", allowing for the reduced replication factor and reduced backend overhead. The tradeoff is a small risk of data inconsistency between the "eventually consistent" Cassandra nodes.

Each instance needs to be assigned to a different VM to avoid an outage caused by the loss of a VM.

If there is enough CPU, memory and disk space, a Cassandra instance can run on the same VM with other stateful service instances of different types. For example, Cassandra can run on the same VM as ZooKeeper.

A separate, dedicated drive is recommended for Cassandra to ensure it receives sufficient IO. The remaining stateful services can share a drive. Any storage option that provides high bandwidth, low latency, and sufficient resiliency is acceptable. For example, SAN and iSCSI provide acceptable performance, provided they are not mounted using the systems network like eth0 or ens. Consider the risk that running on shared storage systems presents. If all of your systems are backed by a single SAN, this configuration introduces a single point of failure.

### **ZooKeeper**

ZooKeeper requires 3 instances to form a quorum, each member would have a complete set of the data.

Each instance needs to be assigned to a different VM to avoid an outage caused by the loss of a VM.

If there is enough CPU, memory and disk space a ZooKeeper instance can run on the same VM alongside other statefulset instances of different types. For example, ZooKeeper can run on the same VM as CouchDB.

While Cassandra needs a dedicated drive, ZooKeeper, Kafka, CouchDB, and Datalayer can share a drive.

### **Kafka**

Kafka can operate with 2 instances for high availability. However, running with 3 for more resiliency is recommended.

With 3 nodes and a replication factor of 3, each node has a copy of 100% of the data. If you expand out from there, each node has a decreasing percentage of the data. Since each node is responsible for less data, adding more nodes increases your cluster capacity. When a node goes offline, the remaining nodes pick up the slack.

As Kafka is the largest set of data, the default replication (when deployed with 2 or more Kafka) has been changed from 3 to 2 to reduce the backend overhead of replicating this data. All other data remains at replication of 3 when deployed with 3 or more Kafka.

Kafka data is organized by topics and spread into partitions. It needs a quorum for some topics, so you can lose only one. For example, if you install with 3, you need 2 available, if you install with 2, you need 1 available.

Each instance needs to be assigned to a different VM to avoid an outage caused by the loss of a VM.

If there is enough CPU, memory, and disk space, a Kafka instance can run on the same VM alongside other statefulset instances of different types. For example, Kafka can run on the same VM as Zookeeper.

While Cassandra needs a dedicated drive, ZooKeeper, Kafka, CouchDB, and Datalayer can share a drive.

### **CouchDB**

CouchDB can operate with 2 instances for high availability. However, running with 3 for more resiliency is recommended. For CouchDB, Cloud Event Management uses the default sharding value of eight shards and 3 replicas. This allows up to eight nodes. You can modify these settings by changing `numShards` and `numReplicas` in the `values.yaml` file. The `numReplicas` parameter controls the replication factor of the data in CouchDB. Scaling CouchDB with `ibm-cem.couchdb.clusterSize` to more nodes than `numReplicas` provides additional capacity and scale.

The default replication factor is already set to 3 even in a deployment with a single CouchDB node.

Each instance needs to be assigned to a different VM to avoid an outage caused by the loss of a VM.

If there is enough CPU, memory and disk space a CouchDB instance can run on the same VM alongside other statefulset instances of different types. For example, CouchDB can run on the same VM as Zookeeper.

While Cassandra needs a dedicated drive, Zookeeper, Kafka, CouchDB, and CEM Data layer can share a drive.

## Datalayer

Datalayer can operate with 2 instances for high availability. However, running with 3 is recommend for more resiliency.

Each instance needs to be assigned to a different VM to avoid an outage caused by the loss of a VM.

If there is enough CPU, memory and disk space a Datalayer instance can run on the same VM alongside other statefulset instances of different types. For example, Datalayer can run on the same VM as Zookeeper.

While Cassandra needs a dedicated drive, ZooKeeper, Kafka, CouchDB, and Datalayer can share a drive.

## Elasticsearch

Elasticsearch can operate with 2 instances for high availability. However, running with 3 is recommend for more resiliency.

Each instance needs to be assigned to a different VM to avoid an outage caused by the loss of a VM.

If there is enough CPU, memory and disk space a Elasticsearch instance can run on the same VM alongside other statefulset instances of different types. For example, Elasticsearch can run on the same VM as Zookeeper.

While Cassandra needs a dedicated drive, ZooKeeper, Kafka, CouchDB, Datalayer, and Elasticsearch can share a drive.

## What are the steps required to set up a highly available environment during a fresh install of the Cloud App Management server?

Setup for a highly available environment is the same as for a non-high availability environment. You will follow the [“Installing IBM Cloud App Management stand-alone on IBM Cloud Private” on page 163](#) procedure, but there are some flags that you must specify in the `prepare-pv.sh` and `pre-install.sh` scripts.

1. Review hardware considerations. The recommended hardware requirements are shown in [Table 9 on page 78](#) in the [“Planning hardware and sizing” on page 77](#) topic.
2. Proceed with your installation as described in [“Installing IBM Cloud App Management stand-alone on IBM Cloud Private” on page 163](#). At the following points, during your installation, you must specify parameters based on your resiliency planning:
  - a. In step 9 in [“Installing IBM Cloud App Management stand-alone on IBM Cloud Private” on page 163](#), run the `prepare-pv.sh` script to prepare persistent volumes, and set the following parameters to match your resiliency plan.
    - Specify **Size1**.
    - For the following parameters, ensure that you list your nodes within quotation marks separated by spaces: **--CassandraNodes**, **--ZookeeperNodes**, **--KafkaNodes**, **--CouchDBNodes**, **--DatalayerNodes**
    - Specify the size in the **--CassandraSize** parameter. For Cassandra, the data is spread evenly between replicas of the statefulset. Each PV needs to be the same size.
    - If IBM Cloud Private uses the IP address instead of the hostname, the IP address is needed.

This example is appropriate for up to 1,000,000 metrics per minute with 3 Cassandra nodes

```
ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh --size1 \  
--cassandraNode "worker13 worker14 worker15" \  
--zookeeperNode "worker1 worker2 worker3" \  
--kafkaNode "worker4 worker5 worker6" \  
--couchdbNode "worker7 worker8 worker9" \  
--datalayerNode "worker10 worker11 worker12"
```

This example is appropriate for over 1,000,000 metrics per minute, with 6 Cassandra nodes. To calculate the metrics needed, see the [How to determine what size to use](#) section.

```
ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh --size1 \  
--cassandraNode "worker01 worker02 worker03 worker07 worker08 worker09" \  
--zookeeperNode "worker04 worker05 worker06" \  
--kafkaNode "worker04 worker05 worker06" \  
--couchdbNode "worker04 worker05 worker06" \  
--datalayerNode "worker04 worker05 worker06"
```

- b. Run the **pre-install.sh** script, set the following parameters to match your resiliency plan.
- Stateful services: For the following parameters, set the replication factor for each statefulset service: **--kafkaClusterSize**, **--zookeeperClusterSize**, **--couchdbClusterSize**, **--datalayerClusterSize**, **--cassandraClusterSize**
  - Stateless services: For the following parameters, set the global replication factor for stateless services: **--minReplicasHPAs**, **--maxReplicasHPAs**
- For more information, see [“Scaling stateless and stateful services”](#) on page 83.

For example:

```
ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/pre-install.sh --accept \  
--releasename ibmcloudappmgmt \  
--namespace default \  
--masterip sample4000.rtp.raleigh.ibm.com \  
--proxyip 9.37.204.30 \  
--proxyhostname sample4000.rtp.raleigh.ibm.com \  
--clusteradomain samplecluster.icp \  
--advanced --cassandraclustersize 3 \  
--kafkaclustersize 3 \  
--zookeeperclustersize 3 \  
--couchdbclustersize 3 \  
--datalayerclustersize 3 \  
--minreplicashpas 2 \  
--maxreplicashpas 3
```

3. Proceed with your installation. If the stateful services do not start, in particular Cassandra, you might need to free up resources, this is described in the [“Creating your service instance”](#) on page 174 topic and also in the [“Freeing up resources for large Pod scheduling”](#) on page 1404 topic.

## Configuring certificates for HTTPS communications

To enable communication between the IBM Cloud App Management server, browsers, and agents, you can configure default or custom certificates.

### Configuring a default certificate

The HTTPS protocol allows communication between the Cloud App Management server and the agents, the server allows connections from the resources that authenticate themselves with a valid certificate. You can configure HTTPS communication that is based on default certificates, which are generated during the installation of the Cloud App Management server.

### Configuring a custom certificate

You might want to use your own certificate instead of the certificate that is generated by the Cloud App Management server. For more information, see [“Configuring a custom server certificate”](#) on page 153.

**Note:** If you do provide your own certificate for the communication between the IBM Cloud App Management server, browsers, and agents, you must create a ConfigMap containing the certificate authority's certificate in x509 PEM format. For example:

```
kubectl create configmap master-ca --from-file=./ca.pem.
```

In this example, the value to enter is master-ca (replace "master-ca" with the name you have chosen).

## Configuring a custom server certificate

Learn how to use your own server certificate instead of the default certificate that is generated by the Cloud App Management server installation. Create certificates, a server key, and a tls secret before you install the Cloud App Management server.

### Before you begin

If you use the HTTPS protocol to communicate between the Cloud App Management server and the agents, the server allows connections from the resources that authenticate themselves with a valid certificate. To enable communication between the IBM® Cloud App Management server, browsers, and agents, you can configure custom certificates.

**Note:** Many certificate authorities have multiple layers of certificates, such as a root certificate and an issuer (or signer) certificate. The `ca.crt` file must be the full chain certificate file. You can use `openssl` to merge certificates. The following example merges the `signer.crt` and `root.crt` files with the `ca.crt` file:

```
openssl x509 -in signer.crt -subject -issuer > ca.crt
openssl x509 -in root.crt -subject -issuer >> ca.crt
```

You must obtain the following files:

- A certificate authority (CA) certificate (`ca.crt`) file that contains the chain of certificates up to (but excluding) the server certificate, with the additional labels that are added during execution of the **openssl x509** commands.
- A server certificate (`server.crt`) file that contains the single certificate used by the server.
- A server private key (`server.key`) file that contains the single private key used by the server.

### About this task

When the Cloud App Management server is installed, a set of signed certificates are created, which are used by the server and agents. You can use your own self-signed certificates or certificates issued by a CA, based on your local security requirements.

You can change the certificate installed on the Cloud App Management server, the agents, or both.

**Note:** Changing the certificate on either the server or the agent causes an interruption in service for all previously connected agents and data collectors. After configuring a custom certificate, you must reconfigure all agents and data collectors to connect to the server. For more information see the following topics:

- [“Configuring the downloaded images” on page 194](#)
- [“Connecting IBM Tivoli Monitoring agents to Cloud App Management server” on page 676](#)
- [“Connecting Cloud APM agents to Cloud App Management server” on page 690](#)
- [Chapter 15, “Deploying ICAM Data Collectors,” on page 555](#)

### Procedure

1. Identify the Kubernetes tls secret files used by IBM Cloud App Management with the following command:

```
kubectl get secret -l release=my_release name
```

Where `my_release name` is the name of your Cloud App Management Helm chart, such as `ibmcloudappgmt`. By default, the secret names are `my_release name-ingress-client` and `my_release name-ingress.tls`.

2. To update the server certificate, complete the following steps:
  - a) Back up the current secret by running the following command:

```
kubectl get secret ibmcloudappgmt-ingress-tls --namespace=my_namespace
```

```
-o yaml > ibmcloudappmgmt-ingress-tls.backup.yaml
```

where `--` denotes an optional parameter and `my_namespace` is the namespace that the installation image file is loaded to.

b) Replace the current secret with your new certificate by running the following command:

```
kubectl create secret generic ibmcloudappmgmt-ingress-tls --namespace=my_namespace
--dry-run -o yaml --from-file=tls.crt=server.crt --from-file=tls.key=server.key
--from-file=ca.crt=ca.crt | kubectl apply -f -
```

c) Optional: If you want to restore your original secret, complete the following steps:

- 1) Open the `my_backup_file.yaml` file with a text editor, such as `vi`, where `my_backup_file` is your backup file name.
- 2) Remove four lines of code from the metadata section, such as the following example:

```
creationTimestamp: 2018-12-04T22:46:57Z
resourceVersion: "6698199"
selfLink: /api/v1/namespaces/default/secrets/ibmcloudappmgmt-ingress-tls
uid: 8122c479-f816-11e8-bb90-00000a150578
```

3) Run the following command:

```
kubectl replace -f my_backup_file
```

3. After configuring a custom certificate, you must redeploy pods by restarting or deleting `cem-users` and `amui` pods. You must also generate the new agent keystore databases by restarting the `agentbootstrap` microservice. Run the following command:

```
kubectl scale --replicas=0 --namespace=my_namespace
deployment ibmcloudappmgmt-agentbootstrap
ibmcloudappmgmt-amui ibmcloudappmgmt-ibm-cem-cem-users
kubectl scale --replicas=1 --namespace=my_namespace
deployment ibmcloudappmgmt-agentbootstrap ibmcloudappmgmt-amui
ibmcloudappmgmt-ibm-cem-cem-users
```

4. If your server custom certificate uses a range scaling algorithm (RSA) key, you must update the agent configuration. Edit the `KDEBE_FIPS_MODE_ENABLED` setting in the `dst_images_dir/global.environment` file, where `dst_images_dir` is the directory to output the configured agent images. If not specified, the configured agent images are saved in the `/depot` folder within the parent directory that contains the agent configuration pack. Change the value from `KDEBE_FIPS_MODE_ENABLED=SuiteB-128` to `KDEBE_FIPS_MODE_ENABLED=SP800-131a`.

### Configuring a custom agent certificate

Learn how to use your own agent certificate instead of the default certificate that is generated by the Cloud App Management server installation. Create certificates, a `tls` secret, and a client key secret before you install the Cloud App Management server.

#### Before you begin

**Note:** Many certificate authorities have multiple layers of certificates, such as a root certificate and an issuer (or signer) certificate. The `ca.crt` file must be the full chain certificate file. You can use `openssl` to merge certificates. The following example merges the `signer.crt` and `root.crt` files with the `ca.crt` file:

```
openssl x509 -in signer.crt -subject -issuer > ca.crt
openssl x509 -in root.crt -subject -issuer >> ca.crt
```

You must obtain the following files:

- A certificate authority (CA) certificate (`ca.crt`) file that contains the chain of certificates up to (but excluding) the agent certificate, with the additional labels that are added during execution of the **openssl x509** commands.
- A client certificate (`client.crt`) file that contains the single certificate used by the agents.

- A client private key (`client.key`) file that contains the single private key used by the agents.

### About this task

When the Cloud App Management server is installed, a set of signed certificates are created, which are used by the server and agents. You can use your own self-signed certificates or certificates issued by a CA, based on your local security requirements.

You can change the certificate installed on the Cloud App Management server, the agents, or both.

**Note:** Changing the certificate on either the server or the agent causes an interruption in service for all previously connected agents and data collectors. After configuring a custom certificate, you must reconfigure all agents and data collectors to connect to the server. For more information see the following topics:

- [“Configuring the downloaded images” on page 194](#)
- [“Connecting IBM Tivoli Monitoring agents to Cloud App Management server” on page 676](#)
- [“Connecting Cloud APM agents to Cloud App Management server” on page 690](#)
- [Chapter 15, “Deploying ICAM Data Collectors,” on page 555](#)

### Procedure

1. Identify the Kubernetes tls secret files used by IBM Cloud App Management with the following command:

```
kubectl get secret -l release=my_release name
```

Where *my\_release name* is the name of your Cloud App Management Helm chart, such as `ibmcloudappgmt`. By default, the secret names are *my\_release name*-`ingress-client` and *my\_release name*-`ingress.tls`.

2. To update the agent certificate, complete the following steps:
  - a) Create a new password file, which protects the keystore databases on the agent machine, by running the following command:

```
echo "password" > client.pass
```

- b) Back up the current secret by running the following command:

```
kubectl get secret ibmcloudappgmt-ingress-client --namespace=my_namespace -o yaml > ibmcloudappgmt-ingress-client.backup.yaml
```

- c) Replace the current secret with your new certificate by running the following command:

```
kubectl create secret generic ibmcloudappgmt-ingress-client --namespace=my_namespace --dry-run -o yaml --from-file=client.crt=client.crt --from-file=client.key=client.key --from-file=ca.crt=ca.crt --from-file=client.password=client.pass | kubectl apply -f -
```

- d) Optional: If you want to restore your original secret, complete the following steps:
  - 1) Open the *my\_backup\_file*.yaml file with a text editor, such as `vi`, where *my\_backup\_file* is your backup file name.
  - 2) Remove four lines of code from the metadata section, such as the following example:

```
creationTimestamp: 2018-12-04T22:46:57Z
resourceVersion: "6698199"
selflink: /api/v1/namespaces/default/secrets/ibmcloudappgmt-ingress-tls
uid: 8122c479-f816-11e8-bb90-00000a150578
```

- 3) Run the following command:

```
kubectl replace -f my_backup_file
```

- After configuring a custom certificate, you must redeploy pods by restarting or deleting `cem-users` and `apmui` pods. You must also generate the new agent keystore databases by restarting the `agentbootstrap` microservice. Run the following command:

```
kubectl scale --replicas=0 --namespace=my_namespace
deployment ibmcloudappgmt-agentbootstrap
ibmcloudappgmt-amui ibmcloudappgmt-ibm-cem-cem-users
kubectl scale --replicas=1 --namespace=my_namespace
deployment ibmcloudappgmt-agentbootstrap ibmcloudappgmt-amui
ibmcloudappgmt-ibm-cem-cem-users
```

- If your server or agent custom certificates use a range scaling algorithm (RSA) key, you must update the agent configuration. Edit the `KDEBE_FIPS_MODE_ENABLED` setting in the `dst_images_dir/global.environment` file, where `dst_images_dir` is the directory to output the configured agent images. If not specified, the configured agent images are saved in the `/depot` folder within the parent directory that contains the agent configuration pack. Change the value from `KDEBE_FIPS_MODE_ENABLED=SuiteB-128` to `KDEBE_FIPS_MODE_ENABLED=SP800-131a`.

## Optimizing performance

Cassandra is the main data store for Cloud App Management. There are some ways to optimize Cassandra performance. You can also configure your drives for server components to improve performance.

### Optimizing disk performance for Cassandra

Best practice for optimizing disk performance for the Cassandra database is to lower the default disk readahead for the drive or partition where your Cassandra data is stored. By default, the Linux kernel reads additional file data so that subsequent reads can be satisfied from the cache. The file access patterns of Cassandra queries result in the readaheads mostly being unused, therefore polluting the cache, driving up I/O time and also results in excessive disk I/O levels.

### Before you begin

You can view your current readahead settings with either of these commands:

```
lsblk --output NAME,KNAME,TYPE,MAJ:MIN,FSTYPE,SIZE,RA,MOUNTPOINT,LABEL
```

```
blockdev --report
```

Examples:

```
lsblk --output NAME,KNAME,TYPE,MAJ:MIN,FSTYPE,SIZE,RA,MOUNTPOINT,LABEL
NAME          KNAME TYPE MAJ:MIN FSTYPE      SIZE  RA MOUNTPOINT LABEL
fd0           fd0  disk   2:0      xfs          4K    128 /
sda           sda  disk   8:0      xfs          80G   4096 /
├─sda1        sda1 part   8:1      xfs          1G    4096 /boot
└─sda2        sda2 part   8:2      LVM2_member 79G   4096 /
   └─rhel-root dm-0  lvm    253:0     xfs          75G   4096 /
      └─rhel-swap dm-1  lvm    253:1     swap         4G    4096 [SWAP]
sdb           sdb  disk   8:16     xfs          100G  4096 /docker
sdc           sdc  disk   8:32     xfs          2T    128
```

```
blockdev --report
RO  RA  SSZ  BSZ  StartSec      Size  Device
rw  256  512  4096  0             4096  /dev/fd0
rw  8192 512  4096  0             85899345920 /dev/sda
rw  8192 512  512   2048         1073741824 /dev/sda1
rw  8192 512  4096  2099200     8482455520 /dev/sda2
rw  8192 512  512  0           107374182400 /dev/sdb
rw  8192 512  512  0           80530636800 /dev/dm-0
rw  8192 512  4096  0           4290772992 /dev/dm-1
rw  256  512  4096  0           2148557389824 /dev/sdc
```

Looking at the `RA` and `Size` columns, the readahead of 8192 combined with the size of 512 results in a readahead of 4096 KB. That means any read on the `/` root drive results in 4 MB of disk I/O into the system cache. Best practice is to use a separate drive for the Cassandra data, as well as the other `StatefulSet` services requiring disk space.



## About this task

The following steps provide two examples of how to modify the readahead settings for an existing drive by using the command line or the tuned service to ensure better performance for Cassandra. You can manually set the readahead on an existing drive or modify the `tuned.service` disk settings to make the readahead settings persistent. These steps need to be performed on each VM running Cassandra.

## Procedure

Manually set readahead on an existing drive or volume.

- To set the readahead, use the `blockdev` command with the internal kernel device name (KNAME):
  - To find the KNAME of the device to modify, run the following command:

```
lsblk --output NAME,KNAME,TYPE,MAJ:MIN,FSTYPE,SIZE,RA,MOUNTPOINT,LABEL
```

In this example, Cassandra is running on `dm-2`. Your KNAME might be different based on your system settings:

NAME	KNAME	TYPE	MAJ:MIN	FSTYPE	SIZE	RA	MOUNTPOINT	LABEL
fd0	fd0	disk	2:0		4K	128		
sda	sda	disk	8:0		80G	4096		
├sda1	sda1	part	8:1	xfs	1G	4096	/boot	
└sda2	sda2	part	8:2	LVM2_member	79G	4096		
├rhel-root	dm-0	lvm	253:0	xfs	75G	4096	/	
└rhel-swap	dm-1	lvm	253:1	swap	4G	4096	[SWAP]	
sdb	sdb	disk	8:16	xfs	100G	4096	/docker	
sdc	sdc	disk	8:32	LVM2_member	2T	4096		
└vg_sdc-lv_cassandra	dm-2	lvm	253:2	xfs	2T	4096	/k8s/data/cassandra	cassandra

- Enter the `blockdev` command with `--setra` in number of blocks (for example, a readahead of 16 with size of 512 bytes results in an 8KB readahead):

```
blockdev --setra 16 device
```

Example: `blockdev --setra 16 /dev/dm-2`

- Verify the readahead settings:

```
lsblk --output NAME,KNAME,TYPE,MAJ:MIN,FSTYPE,SIZE,RA,MOUNTPOINT,LABEL
```

In this example, Cassandra is running on `dm-2`:

NAME	KNAME	TYPE	MAJ:MIN	FSTYPE	SIZE	RA	MOUNTPOINT	LABEL
fd0	fd0	disk	2:0		4K	128		
sda	sda	disk	8:0		80G	4096		
├sda1	sda1	part	8:1	xfs	1G	4096	/boot	
└sda2	sda2	part	8:2	LVM2_member	79G	4096		
├rhel-root	dm-0	lvm	253:0	xfs	75G	4096	/	
└rhel-swap	dm-1	lvm	253:1	swap	4G	4096	[SWAP]	
sdb	sdb	disk	8:16	xfs	100G	4096	/docker	
sdc	sdc	disk	8:32	LVM2_member	2T	4096		
└vg_sdc-lv_cassandra	dm-2	lvm	253:2	xfs	2T	8	/k8s/data/cassandra	cassandra

- If you are modifying a running environment, restart the Cassandra Docker container to use the new readahead values.

You can restart the Cassandra Docker container as an IBM Cloud Private admin, either through the IBM Cloud Private UI or `kubectl`.

With this method, the `tuned` service adjusts the configuration settings to optimize system performance. Tuned profiles overwrite the smaller readahead setting used in the LVM setup. The `tuned` service adjusts the configuration settings to optimize system performance. The service can modify settings such as disk device readahead. Tuned profiles overwrite the smaller readahead setting used in the LVM setup. To prevent the overwrite, add the setting to your tuned profile. For more information, see [Performance tuning with tuned and tuned-adm](#) in the *Red Hat Performance Tuning Guide*.

- Modify `tuned.service` disk settings to make readahead persistent:
  - Format a blank drive for the Cassandra data to be stored as described in [“Configuring the disk drives for services”](#) on page 158.

- b) Use the **tuned-adm** command to see the current active profile:

```
tuned-adm active
Current active profile: virtual-guest
```

The output in this example shows that the active profile is `virtual-guest`. Note: Your profile and configuration may be different.

- c) Copy the profile to the `/etc/tuned` directory.

The default profile definitions are stored in `/usr/lib/tuned/`.

In our example, the definitions are in `/usr/lib/tuned/virtual-guest/tuned.conf`. The definitions for `virtual-guest` contain **`include=throughput-performance`**, which means the settings inherit the settings of `throughput-performance`. Looking at `/usr/lib/tuned/throughput-performance/tuned.conf`, we see that this is where the **`readahead=>4096`** is being set.

```
cp -a /usr/lib/tuned/throughput-performance/ /etc/tuned/
```

- d) Copy the profile to the `/etc/tuned` directory.

The default profile definitions are stored in `/usr/lib/tuned/`.

In our example, the definitions are in `/usr/lib/tuned/virtual-guest/tuned.conf`. The definitions for `virtual-guest` contain **`include=throughput-performance`**, which means the settings inherit the settings of `throughput-performance`. Looking at `/usr/lib/tuned/throughput-performance/tuned.conf`, we see that this is where the **`readahead=>4096`** is being set.

```
cp -a /usr/lib/tuned/throughput-performance/ /etc/tuned/
```

- e) Add the following section to the `/etc/tuned/throughput-performance/tuned.conf` file, making sure that it is above the existing `[disk]` section.

```
[disk-cassandra]
type=disk
devices=dm-2
readahead=8
```

- f) Reload the tuned profile:

```
tuned-adm profile virtual-guest
```

- g) Verify the new `readahead` setting:

```
lsblk --output NAME,KNAME,TYPE,MAJ:MIN,FSTYPE,SIZE,RA,MOUNTPOINT,LABEL
```

Cassandra is running on `dm-2` in this example:

NAME	KNAME	TYPE	MAJ:MIN	FSTYPE	SIZE	RA	MOUNTPOINT	LABEL
fd0	fd0	disk	2:0		4K	128		
sda	sda	disk	8:0		80G	4096		
├sda1	sda1	part	8:1	xfs	1G	4096	/boot	
└sda2	sda2	part	8:2	LVM2_member	79G	4096	/	
├rhel-root	dm-0	lvm	253:0	xfs	75G	4096	/	
└rhel-swap	dm-1	lvm	253:1	swap	4G	4096	[SWAP]	
sdb	sdb	disk	8:16	xfs	100G	4096	/docker	
sdc	sdc	disk	8:32	LVM2_member	2T	4096		
└vg_sdc-lv_cassandra	dm-2	lvm	253:2	xfs	2T	8	/k8s/data/cassandra	cassandra

- h) If you are modifying a running environment, restart the Cassandra Docker container to use the new `readahead` values.

You can restart the Cassandra Docker container as an IBM Cloud Private admin, either through the IBM Cloud Private UI or `kubectl`.

### Configuring the disk drives for services

Set up the drives that are required for your Cloud App Management server components. IBM Cloud App Management requires 6 persistent volumes. For performance and scalability, we recommend using local storage. The steps below provide examples of how to format drives and partition them for use. It is

recommended to use a separate drive for Cassandra. This drive will handle the majority of the disk IO, as well as require separate tuning to optimize IO (see Disk Performance Optimization For Cassandra). For this example, our system has been provisioned with a 2TB drive `/dev/sdc` for Cassandra and a 500GB drive `/dev/sdd` for the other 5 services (Zookeeper, Kafka, CouchDB, Datalayer and Elasticsearch) volumes.

## Procedure

Complete these steps to format the blank drives using logical volumes:

1. Identify the disk: `fdisk -l`

The output in this example shows that the system has been provisioned with a 2000 GB `/dev/sdc` for Cassandra and a 500 GB `/dev/sdd` for the other services:

```
fdisk -l
...
Disk /dev/sdc: 2148.6 GB, 2148557389824 bytes, 4196401152 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk /dev/sdd: 536.9 GB, 536870912000 bytes, 1048576000 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

2. Create physical volumes for each drive: `pvcreate path_to_new_volume -f`

In this example, the commands create 2 new physical volume `/sdc` and `/sdd` as subdirectories of `/dev`:

```
pvcreate /dev/sdc -f
Physical volume "/dev/sdc" successfully created.
```

```
pvcreate /dev/sdd -f
Physical volume "/dev/sdd" successfully created.
```

3. Create a volume group for each physical volume: `vgcreate vg_name pv_path`

In this example, the command creates volume groups `vg_sdc` for the physical volume `/dev/sdc` and `vg_sdd` for the physical volume `/dev/sdd`:

```
vgcreate vg_sdc /dev/sdc
Volume group "vg_sdc" successfully created
vgcreate vg_sdd /dev/sdd
Volume group "vg_sdd" successfully created
```

4. Create a logical volume for Cassandra on volume group `vg_sdc`. Note the volume group used:

```
lvcreate --name lv_name --size 1999G vg_name -y --readahead 8
```

In this example, the command creates a logical volume named `lv_cassandra`, sized at 1999G in the volume group named `vg_sdc`. Note that using size 2000G could result in error, Volume group "vg\_sdc" has insufficient free space (511999 extents): 512000 required.

```
lvcreate --name lv_cassandra --size 1999G vg_sdc -y --readahead 8
Logical volume "lv_cassandra" created.
```

Create the other logical volumes for the other 5 services using the `lvcreate` command on the volume group `vg_sdd`:

```
lvcreate --name lv_kafka --size 200G vg_sdd -y
```

```
lvcreate --name lv_zookeeper --size 1G vg_sdd -y
```

```
lvcreate --name lv_couchdb --size 50G vg_sdd -y
```

```
lvcreate --name lv_datalayer --size 5G vg_sdd -y
```

```
lvcreate --name lv_elasticsearch --size 1G vg_sdd -y
```

5. Format the new logical volumes using the XFS format:

In this example, the command formats the `lv_cassandra` logical volume in the `/dev/vg_sdc/` volume group in XFS format:

```
mkfs.xfs -L cassandra /dev/vg_sdc/lv_cassandra
```

```
meta-data=/dev/vg_sdc/lv_cassandra isize=512    agcount=4, agsize=131006464 blks
        =                               sectsz=512    attr=2, projid32bit=1
        =                               crc=1        finobt=0, sparse=0
data      =                               bsize=4096   blocks=524025856, imaxpct=5
        =                               sunit=0      swidth=0 blks
naming    =version 2                       bsize=4096   ascii-ci=0 ftype=1
log       =internal log                   bsize=4096   blocks=255872, version=2
        =                               sectsz=512   sunit=0 blks, lazy-count=1
realtime  =none                           extsz=4096   blocks=0, rtextents=0
```

Format the remaining logical volumes for each of the services:

```
mkfs.xfs -L kafka /dev/vg_sdd/lv_kafka
```

```
mkfs.xfs -L zk /dev/vg_sdd/lv_zookeeper
```

```
mkfs.xfs -L couchdb /dev/vg_sdd/lv_couchdb
```

```
mkfs.xfs -L datal /dev/vg_sdd/lv_datalayer
```

```
mkfs.xfs -L elastic /dev/vg_sdd/lv_elasticsearch
```

6. Create the directories for each filesystem:

```
mkdir -p /k8s/data/cassandra
```

```
mkdir -p /k8s/data/kafka
```

```
mkdir -p /k8s/data/zookeeper
```

```
mkdir -p /k8s/data/datalayer
```

```
mkdir -p /k8s/data/couchdb
```

```
mkdir -p /k8s/data/elasticsearch
```

7. Add the new filesystem directories to `/etc/fstab`

```
echo "/dev/vg_sdc/lv_cassandra /k8s/data/cassandra xfs defaults 0 0" >> /etc/fstab
```

```
echo "/dev/vg_sdd/lv_kafka /k8s/data/kafka xfs defaults 0 0" >> /etc/fstab
```

```
echo "/dev/vg_sdd/lv_zookeeper /k8s/data/zookeeper xfs defaults 0 0" >> /etc/fstab
```

```
echo "/dev/vg_sdd/lv_datalayer /k8s/data/datalayer xfs defaults 0 0" >> /etc/fstab
```

```
echo "/dev/vg_sdd/lv_couchdb /k8s/data/couchdb xfs defaults 0 0" >> /etc/fstab
```

```
echo "/dev/vg_sdd/lv_elasticsearch /k8s/data/elasticsearch xfs defaults
0 0" >> /etc/fstab
```

8. Mount the new filesystems on the new directories:

```
mount /k8s/data/cassandra
```

```
mount /k8s/data/kafka
```

```
mount /k8s/data/zookeeper
```

```
mount /k8s/data/datalayer
```

```
mount /k8s/data/couchdb
```

```
mount /k8s/data/elasticsearch
```

9. Verify the mount point and readahead settings with the `lsblk` command: **`lsblk --output NAME, KNAME, TYPE, MAJ:MIN, FSTYPE, SIZE, RA, MOUNTPOINT, LABEL`**

In this example, the characteristics of the mount points are displayed:

```
lsblk --output NAME, KNAME, TYPE, MAJ:MIN, FSTYPE, SIZE, RA, MOUNTPOINT, LABEL
NAME          KNAME  TYPE MAJ:MIN FSTYPE  SIZE  RA MOUNTPOINT  LABEL
fd0           fd0    disk  2:0     4K     128
sda           sda    disk  8:0     80G   4096
├─sda1        sda1   part  8:1     xfs     1G   4096 /boot
├─sda2        sda2   part  8:2     LVM2_member 79G  4096
│ └─rhel-root dm-0   lvm   253:0   xfs     75G  4096 /
│ └─rhel-swap dm-1   lvm   253:1   swap    4G   4096 [SWAP]
sdb           sdb    disk  8:16    xfs    100G  4096 /docker
sdc           sdc    disk  8:32    LVM2_member 2T    128
├─vg_sdc-lv_cassandra dm-2   lvm   253:2   xfs     2T    4 /k8s/data/cassandra
cassandra
sdd           sdd    disk  8:48    LVM2_member 500G  128
├─vg_sdd-lv_kafka dm-3   lvm   253:3   xfs    200G  128 /k8s/data/kafka kafka
├─vg_sdd-lv_zookeeper dm-4   lvm   253:4   xfs     1G    128 /k8s/data/zookeeper zk
├─vg_sdd-lv_couchdb dm-5   lvm   253:5   xfs    50G   128 /k8s/data/couchdb couchdb
├─vg_sdd-lv_datalayer dm-6   lvm   253:6   xfs     5G    128 /k8s/data/datalayer datal
└─vg_sdd-lv_elasticsearch dm-7   lvm   253:7   xfs     1G    128 /k8s/data/elasticsearch
elastic
```

10. To increase the size of a volume, use the **`lvextend`** command, for example:

```
lvextend -L 20G /dev/mapper/vg_sdd-lv_couchdb
```

11. After extending the volume, resize the xfs directory, for example:

```
xfs_growfs /k8s/data/couchdb
```

**Note:** The kubernetes persistent volume definitions' "capacity" are not hard limits. The persistent volumes will use whatever storage is available to them inside the directory. However, after increasing a volume's capacity, it is recommended to modify the persistent volume's capacity for consistency.

### Related tasks

[“Optimizing disk performance for Cassandra” on page 156](#)

## Cassandra Disk and Memory Usage - Adding Memory Resources

### Cassandra kernel file system cache

Cassandra JVMs perform large amounts of disk I/O operations, for example, writing new data, compacting existing SSTables, and reading for queries. Cassandra relies on the kernel file system cache for optimizing reads. Recent and frequently used files are kept in the memory cache. It would be impractical to keep all of the hundreds of GBs of metric data Cassandra stores in memory at scale. Therefore, it is impossible to completely eliminate disk reads. By default, Cassandra containers are given 16 GB of RAM, set in their Kubernetes resource requests and limits. After the JVM, approximately 6 GB remains for file caching. This file cache size satisfies many of the most common

reads from memory, for example: recent metric data or frequently accessed tables like topology or events. Operations like SSTable compaction, are also generally able to be completed from memory. SSTable compaction involves merging the many immutable files Cassandra stores data in into a smaller number of larger files.

### **Metric summarization**

Metric summarization in particular can be heavy on the disk reads. For more information on metric summarization, see [“Configuring summarization”](#) on page 766.

While metric summarization is generally requesting recent data, it is possible for other operations to claim the file cache, for example, SSTable compaction. Given the scale of the metric summarization requests, this hit on reads is larger than normal UI query loads.

### **Cassandra I/O**

Cassandra I/O is unique compared to many traditional databases. SSTable compaction results in large amounts of reads/writes, but the read/writes are generally sequential I/O since it is reading and writing straight through files. This makes it ideal for traditional hard disk drive arrays that do well on sequential operations. However, queries can be much more random in their I/O pattern, for example, searching the headers of SSTables and pulling specific rows of data from the tables. For this reason, it is recommended by the Cassandra community to tune down the read ahead setting. This helps to reduce the wasted I/O reading data that is never used in queries, for more information, see [“Optimizing disk performance for Cassandra”](#) on page 156. The result is that the disks needed to support Cassandra, need the ability to do both large amounts of reads/writes, as well as handle random I/O well. This is why we do not recommend any network based storage solutions, as their performance and ability to handle the bandwidth is generally insufficient.

### **Increase the Cassandra memory request and limit**

In containers like Docker and Kubernetes, the kernel filesystem cache is limited to the space within the container. For example, if you run Cassandra on a system with 64 GB of RAM, but the Cassandra container memory limit is 16 GB, the additional 48 GB of RAM will be unavailable to Cassandra to optimize the disk I/O. In order to take advantage of the additional RAM on the system, the Cassandra memory request and limit need to be increased.

Since this kernel filesystem cache is included in the container resource usage, it is recommended by the Kubernetes community to set the memory request and limit equal for containers that do disk I/O and therefore use the kernel filesystem cache. This is because of the nature of the kernel filesystem cache. Unused RAM is wasted RAM; therefore the kernel tries to use as much memory as it can to optimize disk I/O by keeping the information in memory. Therefore, the container will always eventually almost reach the limit as it files cache. So requesting less than the limit will misrepresent the amount of memory the container will use on a system. Note, the kernel filesystem cache memory (buffers/cache) is still available to be freed and given to things that need it, like processes.

Increasing the memory of your Cassandra containers, and therefore the kernel filesystem cache space available to Cassandra, will allow disk I/O to be reduced as more (but not all) operations will be satisfied from memory. This can be done by running:

1. Run the following command:

```
kubect1 edit
statefulset my_release_name-cassandra
```

2. Modifying both the resource requests and limits. The existing size1 settings are:

```
resources:
  limits:
    cpu: "6"
    memory: 16Gi
  requests:
    cpu: "4"
    memory: 16Gi
```

To double the total memory available and add approximately 16 GB to the kernel filesystem cache, follow this example:

```
resources:
  limits:
    cpu: "6"
    memory: 32Gi
  requests:
    cpu: "4"
    memory: 32Gi
```

Future upgrades or helm updates can reset these values if the yaml is not modified as well. The definitions for the resources are contained in the `_resources.tpl` of the Cloud App Management charts, located here: `ibm-cloud-appmgmt-prod/charts/cassandra/templates/_resources.tpl`. Modify the values in the size corresponding to your environment deployment. For example:

```
size1:
  replicas: 3
  cassandraHeapSize: "8G"
  cassandraHeapNewSize: "2G"
  cassandraConcurrentCompactors: 4
  cassandraMemtableFlushWriters: 2
  resources:
    requests:
      memory: "16Gi"
      cpu: "4"
    limits:
      memory: "16Gi"
      cpu: "6"
```

To double the total memory available and add roughly 16 GB to the kernel file system cache:

```
size1:
  replicas: 3
  cassandraHeapSize: "8G"
  cassandraHeapNewSize: "2G"
  cassandraConcurrentCompactors: 4
  cassandraMemtableFlushWriters: 2
  resources:
    requests:
      memory: "32Gi"
      cpu: "4"
    limits:
      memory: "32Gi"
      cpu: "6"
```

**Note:** This custom modification needs to be done for your charts in all subsequent releases. If the resources are not modified in the charts, the values reset and you need to run `kubectl edit statefulset my_release_name-cassandra`. It is also possible to modify the values during the helm upgrade instead of `kubectl edit`.

## Installing IBM Cloud App Management stand-alone on IBM Cloud Private

Learn how to install the Cloud App Management server with the IBM Cloud Private CLI by running the `helm install` command.

### Before you begin

1. Download and install the IBM Cloud Private V3.2.1 for Linux (x86\_64) Docker package (Part number: CC3KPEN) from [IBM Passport Advantage®](#). Find the installation image by searching for it using its part number. Then, install it's most recent fix pack from [IBM Fix Central](#). Search for IBM Cloud Private to find its associated fix packs.
2. It is best practice to install the Cloud App Management server in a new, non-default namespace. There was a new `limitrange` resource added in IBM Cloud Private V3.2.1 for the default namespace. Installing the Cloud App Management server into the default namespace of IBM Cloud Private V3.2.1

can be impacted by this limit range. If you are not using a shared cluster, you can delete the `limitrange` resource in the default namespace by running the following command:

```
kubectl delete limitrange default-limit -n default
```

3. Review the optimization for performance topics. For more information, see [“Optimizing performance” on page 156](#).
4. If you want to use a custom certificate, you must complete some steps before you install the Cloud App Management server. For more information, see [“Configuring a custom server certificate” on page 153](#).
5. You must onboard LDAP users. For more information, see [“Onboarding LDAP users” on page 94](#).
6. The Cloud App Management server uses the UIDs 100, 1000 and 1001 and GIDs 100, 1000 and 1001. To avoid any issues with file ownership or permissions, you can create users and groups for each UID and GID. For example, create a user name "icam" with UID 100 and a group "icamgrp" with GID 100. Create the users and groups for UID 1000 and 1001 and GID 1000 and 1001. You can choose any user and group names when you create them. While not required, it is considered best practice to create these users and groups before you install to help avoid any confusion with file and process ownership. If any users or groups exist with the UIDs or GIDs 100, 1000, or 1001 then you might observe files and processes that are owned by those users and groups.
7. Configure the Elasticsearch `vm.max_map_count` parameter on all worker nodes by completing the following steps:
  - a. For Elasticsearch, you must set a kernel parameter on all worker nodes. These nodes are identified when you configure the persistent storage later in this procedure. You must set `vm.max_map_count` to a minimum value of 1048575. Set the parameter by using the `sysctl` command to ensure that the change takes effect immediately. Run the following command on each worker node:

```
sysctl -w vm.max_map_count=1048575
```

- b. You must also set the `vm.max_map_count` parameter in the `/etc/sysctl.conf` file to ensure that the change is still in effect after the node is restarted.

```
vm.max_map_count=1048575
```

## Procedure

Complete the following steps as an administrator:

1. Locate and download the Cloud App Management server installation image file `icam_ppa_2019.4.0_prod.tar.gz` (Part number: CC4KNEN) from [IBM Passport Advantage](#). Find the installation image by searching for it using its part number. For more information on the IBM Cloud App Management components and their part numbers, see [“Part numbers” on page 71](#).
2. As the administrator, log in to the management console. Run the following command:

```
cloudctl login -a my_cluster_URL -n my_namespace --skip-ssl-validation
```

Where `my_cluster_URL` is the name that you defined for your cluster such as `https://cluster_address:443` and `my_namespace` is the namespace where you are installing your Cloud App Management server. For future references to `masterIP`, use the value you are using for `cluster_address`. A `cluster_address` address example is: `icp-console.apps.organic-test-icp-mst.domain.com`.

3. Log in to docker registry:

```
docker login my_cluster_CA_domain:8500
```

where `my_cluster_CA_domain` is the certificate authority (CA) domain, such as `mycluster.icp`. If you did not specify a `my_cluster_CA_domain`, the default value is `mycluster.icp`.

4. Create an installation directory. You can specify any name for the directory.



```
mkdir -p install_dir
```

**Note:** The Cloud App Management server installation image (PPA file) is large. Ensure that you have enough free space to store the file.

5. Move the Passport Advantage Archive (PPA) file: `icam_ppa_2019.4.0_prod.tar.gz` into the `install_dir` directory and go to the `install_dir` directory.

```
mv app_mgmt_server_2019.4.0.tar.gz install_dir/  
cd install_dir
```

**Note:** The Cloud App Management server installation image (PPA file) is large. Ensure that you have enough free space to store the file.

6. Extract the Helm chart from the Cloud App Management server installation image file: `icam_ppa_2019.4.0_prod.tar.gz` and decompress it.

```
tar -xvf icam_ppa_2019.4.0_prod.tar.gz charts
```

```
tar -xvf charts/ibm-cloud-appmgmt-prod-1.6.0.tgz
```

7. Load the image archive into the catalog. This process can take 10 - 30 minutes.

```
cloudctl catalog load-archive --archive ./icam_ppa_2019.4.0_prod.tar.gz [--registry  
my_cluster_CA_domain:8500] [--repo my_helm_repo_name]
```

8. Create local directories on the selected IBM common services worker nodes where you want the statefulsets to run. Record the IP address of each worker node where the directories are created. You need these IP addresses when you run the `prepare-pv.sh` script in the next step. For more information about optimization performance, see [“Optimizing performance” on page 156](#).

**Note:** The Cassandra persistent volume should reside on its own disk and dedicated worker node. Avoid creating the other persistent volumes on the same worker node and the same persistent volume with Cassandra.

**Note:** The ZooKeeper and Kafka persistent volumes should reside on the same worker node.

**Note:** The CouchDB and Datalayer persistent volumes should reside on the same worker node.

- Cassandra persistent storage:

```
ssh root@cassandraNode  
mkdir -p /k8s/data/cassandra
```

- ZooKeeper persistent storage:

```
ssh root@zookeeperNode  
mkdir -p /k8s/data/zookeeper
```

- Kafka persistent storage:

```
ssh root@kafkaNode  
mkdir -p /k8s/data/kafka
```

- CouchDB persistent storage:

```
ssh root@couchdbNode  
mkdir -p /k8s/data/couchdb
```

- Datalayer persistent storage:

```
ssh root@datalayerNode  
mkdir -p /k8s/data/datalayer
```

- Elasticsearch persistent storage

```
mkdir -p /k8s/data/elasticsearch
```

9. Create the storage classes and persistent volume yam files for each of the following Cloud App Management statefulsets by running the `ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh` script. The use of persistent volumes that are backed by local disks or local partitions is highly recommended.

**Note:** For more information about storage, after you select the target **ibm-cloud-appmgmt-prod** chart from the catalog, select the **Overview** tab if it is not already selected. Next, scroll down the Storage section on the IBM Cloud App Management page on the right. For more information about how to configure persistent storage, see the following topics: [“Planning hardware and sizing” on page 77](#), [Understanding Kubernetes storage](#), [Planning a storage solution](#), and [Planning persistent storage](#).

```
Usage: ./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh
  --releaseName <name>           Release name (default is ibmcloudappmgmt)
  --size0_amd64                   Install as size0 on amd64 (minimum resource
requirements) - if omitted, specify each size in parameters
  --size0_ppc64le                 Install as size0 on ppc64le (minimum resource
requirements) - if omitted, specify each size in parameters
  --size1_amd64                   Install as size1 on amd64 (standard resource
requirements) - if omitted, specify each size in parameters
  --size1_ppc64le                 Install as size1 on ppc64le (standard resource
requirements) - if omitted, specify each size in parameters

*Required flags for local storage:
  --local                         Use local persistent volume storage
  --CassandraNodes <worker_node_IP> Worker node(s) for Cassandra. For multiple
Cassandra use a quoted, space separated list.
  --KafkaNodes <worker_node_IP>   Worker node(s) for Kafka. For multiple Kafka use
a quoted, space separated list.
  --ZookeeperNodes <worker_node_IP> Worker node(s) for Zookeeper. For multiple
Zookeeper use a quoted, space separated list.
  --CouchDBNodes <worker_node_IP>  Worker node(s) for CouchDB. For multiple CouchDB
use a quoted, space separated list.
  --DatalayerNodes <worker_node_IP> Worker node(s) for Datalayer. For multiple
Datalayer use a quoted, space separated list.
  --ElasticsearchNodes <worker_node_IP> Worker node(s) for Elasticsearch. For multiple
Elasticsearch use a quoted, space separated list.

*Optional storage directory paths for local storage:
  --CassandraDir <directory>      Local system directory for Cassandra (default
is /k8s/data/cassandra)
  --CassandraBackupDir <directory> Local system directory for Cassandra backups
(default is /k8s/data/cassandra_backup)
  --KafkaDir <directory>          Local system directory for Kafka (default is /k8s/
data/kafka)
  --ZookeeperDir <directory>      Local system directory for Zookeeper (default
is /k8s/data/zookeeper)
  --CouchDBDir <directory>        Local system directory for CouchDB (default is /k8s/
data/couchdb)
  --DatalayerDir <directory>      Local system directory for Datalayer (default
is /k8s/data/datalayer)
  --ElasticsearchDir <directory>  Local system directory for Elasticsearch (default
is /k8s/data/elasticsearch)

*Optional storage class name flags for local storage:
  --CassandraClass <className>    Storage class name for Cassandra (default is
<release_name>-local-storage-cassandra)
  --CassandraBackupClass <className> Storage class name for Cassandra backups
(default is <release_name>-local-storage-cassandra-backup)
  --KafkaClass <className>        Storage class name for Kafka (default is
<release_name>-local-storage-kafka)
  --ZookeeperClass <className>    Storage class name for Zookeeper (default is
<release_name>-local-storage-zookeeper)
  --CouchDBClass <className>      Storage class name for CouchDB (default is
<release_name>-local-storage-couchdb)
  --DatalayerClass <className>    Storage class name for Datalayer (default is
<release_name>-local-storage-datalayer)
  --ElasticsearchClass <className> Storage class name for Elasticsearch (default
is <release_name>-local-storage-elasticsearch)

*Required flags for vSphere storage:
  --vSphere                       Use vSphere provisioned storage (requires existing
vSphere storage class)

*Optional storage size flags for local and vSphere storage:
  --CassandraSize <size>          Size of persistent volume for Cassandra
(default size0_amd64 and size0_ppc64le is 50Gi)
```

```

--CassandraBackupSize <size>          Size of persistent volume for Cassandra
backups (default size0_amd64 and size0_ppc64le is 50Gi)
--KafkaSize <size>                    Size of persistent volume for Kafka (default
size0_amd64 and size0_ppc64le is 5Gi)
--ZookeeperSize <size>                Size of persistent volume for Zookeeper
(default size0_amd64 and size0_ppc64le is 1Gi)
--CouchDBSize <size>                  Size of persistent volume for CouchDB
(default size0_amd64 and size0_ppc64le is 5Gi)
--DatalayerSize <size>                 Size of persistent volume for Datalayer
(default size0_amd64 and size0_ppc64le is 5Gi)
--ElasticsearchSize <size>            Size of persistent volume for Elasticsearch
(default size0_amd64 and size0_ppc64le is 5Gi)

```

**Note:** The `--local` and `--vSphere` are incompatible. You must choose one or the other. Local storage is highly recommended. If you don't choose either the `--local` or `--vSphere` options, the `prepare-pv.sh` defaults to local storage.

```
kubectl get nodes
```

You don't need to use all the parameters, here are some example scenarios with the typical parameters used:

Scenario 1: No High Availability with a total of 2 VMs. 1 for Cassandra, 1 for the remaining statefulsets:

```

ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh --size1_amd64
--cassandraNode "worker01" --zookeeperNode "worker02"
--kafkaNode "worker02" --couchdbNode "worker02" --datalayerNode "worker02" --
elasticsearchNode "worker02"

```

Scenario 2: High Availability with a total of 6 VMs. 3 for Cassandra, 3 for the remaining statefulsets:

```

ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh --size1_amd64
--cassandraNode "worker01 worker02 worker03" --zookeeperNode "worker04 worker05 worker06"
--kafkaNode "worker04 worker05 worker06" --couchdbNode "worker04 worker05 worker06"
--datalayerNode "worker04 worker05 worker06" --elasticsearchNode "worker04 worker05
worker06"

```

Scenario 3: High Availability with a total of 9 VMs. 6 for Cassandra, 3 for the remaining statefulsets:

```

ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh
--size1_amd64 --cassandraNode "worker01 worker02 worker03 worker07 worker08 worker09"
--zookeeperNode "worker04 worker05 worker06" --kafkaNode "worker04 worker05 worker06" --
couchdbNode "worker04 worker05 worker06"
--datalayerNode "worker04 worker05 worker06" --elasticsearchNode "worker04 worker05
worker06"

```

- The yml files that make up your storage classes and persistent volumes are now created. Next, you must create the Kubernetes resources that will use them. Navigate to the `ibm-cloud-appmgmt-prod/ibm_cloud_pak/yaml/` directory, and create the storage classes and persistent volumes in there:

```

cd ibm-cloud-appmgmt-prod/ibm_cloud_pak/yaml/
kubectl create -f . -n my_namespace
cd ../../../../ # this should move you back to the install_dir directory

```

- Run the `ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/pre-install.sh` script.

The `pre-install.sh` completes many functions. At a minimum, the script creates a `values.yaml` file that you can use to override the default values set in the Cloud App Management Helm chart. It can also be used to enable TLS encryption between the agents and the Cloud App Management server. Refer to the following code snippet to check which flags are required and optional. Examples are given after the code snippet.

```

Usage ./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/pre-install.sh
Use this script to perform preparation tasks that require admin permissions before IBM
Cloud AppMgmt is installed.

```

```

*Required flags
--accept

```

```

Accepts license agreement(s)

```

```

--https                                Install with HTTPS enabled (HTTPS is always
enabled in Advanced offering)
--advanced                              Install as ADVANCED offering (omit this
parameter will install as Base offering)
--releaseName <name>                   Release name (default is ibmcloudappmgmt)
--masterAddress <FQDN>                  Fully qualified domain name (FQDN) for the ICP
Master.
                                        In a highly available environment, this would be
the FQDN of the HAProxy or load balancer for ICP.
--masterPort <int>                       The port of the ICP master. On OpenShift the
default port is 443, which will need to be specified here. (default is 8443)
--ingressPort <int>                       The ingress port used to access the ICP console.
(default is 443)
--proxyIP <IP>                           IP address for ICP Proxy.
                                        In a highly available environment, this would be
the IP address of the HAProxy or load balancer for ICP.
--proxyFQDN <FQDN>                       Fully qualified domain name (FQDN) for the ICP
Proxy.
                                        In a highly available environment, this would be
the FQDN of the HAProxy or load balancer for ICP.
--namespace <name>                       Namespace (default is default)
--clusterCAdomain <name>                 ICP cluster domain name, default is mycluster.icp
                                        This value, combined with the values provided
for --repositoryPort and --namespace, will determine the imageRepository where ICAM will
look for images, e.g. with --clusterCAdomain mycluster.icp --repositoryPort 8500 and --
namespace icam, you will end up with an imageRepository of mycluster.icp:8500/icam
--repositoryPort <int>                   The port of the image repository. On OpenShift
the default port is 5000, which will need to be specified here. (default is 8500)

*Optional - Email setup:
--emailtype <smtp|api>                   Type of email, either smtp or api
--emailfrom <emailAddress>               Email address to show on sent mail as from
--smtpHost <hostname>                     SMTP hostname
--smtpport <port>                         SMTP port
--smtpuser <user>                         SMTP user
--smtpPass <password>                     SMTP password
--smtpauth <true|false>                  User authentication required for SMTP connection
(default is true)
--smtprejectUnauthorized <true|false>    Set this to false to allow self signed
certificates when connecting via TLS, true enforces TLS authorization checking (default is
true)
--apikey <key>                            API key file

*Optional - High availability and horizontal scale settings
--minReplicasHPAs <int>                   The minimum number of replicas for each
deployment, controlled by HPAs
--maxReplicasHPAs <int>                   The maximum number of replicas for each
deployment, controlled by HPAs
--kafkaClusterSize <int>                  The number of Kafka replicas (the replication
factor for Kafka topics will be set to this value, up to a max of 3)
--zookeeperClusterSize <int>              The number of Zookeeper replicas (all Zookeeper
data is replicated to all zookeeper nodes)
--couchdbClusterSize <int>                 The number of CouchDB replicas (the CouchDB data
data replication defaults to 3, even if the cluster has 1 or 2 nodes)
--datalayerClusterSize <int>              The number of Datalayer replicas (the datalayer
relies on Kafka and internal jobs for handling data replication)
--elasticsearchClusterSize <int>          The number of Elasticsearch replicas (the number
of replica shards is determined from the number of Elasticsearch instances)
--redisServerReplicas <int>               The number of redis server replicas. Defaults to
1
--cassandraClusterSize <int>              The number of Cassandra replicas (the
replication factor for Cassandra keyspaces will be set to this value, up to a max of 3)
--cassandraUsername <string>              The username Cassandra will use. You must use a
username other than 'cassandra'. If left unset, the default cassandra credentials will be
used.

*Optional - Other
--metricSummarization <string>            Enables or disables metric summarization. Set to
'true' or 'false'. Defaults to 'false' if not specified.
--metricC8Rep <replication_string>        The replication string for the metric data
(default is '{"class':'SimpleStrategy','replication_factor':X}", where X is the
cassandraClusterSize up to 2)
--openttC8Rep <int>                       The replication factor for the Open Transaction
Tracking data (default is to match cassandraClusterSize up to 2)
--metricKafkaRep <int>                    The replication factor for the metric Kafka data
(default is to match kafkaClusterSize up to 2)
--useTLSCertsJob <true|false>             Enables or disables creating Ingress TLS
Certificates using Kubernetes Job. Set to 'true' or 'false'. Defaults to 'false' if not
specified.

Example of install an Advanced offering with HTTPS enabled on Openshift:
./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/pre-install.sh --accept \

```

```

--releasename icam \
--namespace default \
--masteraddress x.xx.xx.xx \
--proxyip x.xx.xx.xx \
--proxyfqdn proxy.example.com \
--clustercadomain docker-registry.default.svc \
--repositoryPort 5000 \
--advanced \
--cassandraUsername customCassandraSuperuser \
--masterPort 443 \
--ingressPort 443

```

Example of install an Advanced offering with HTTPS enabled on Openshift, using high availability:

```

./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/pre-install.sh --accept \
--releasename icam \
--namespace default \
--masteraddress haproxy.example.com \
--proxyip x.xx.xx.xx \
--proxyfqdn haproxy.example.com \
--clustercadomain docker-registry.default.svc \
--repositoryPort 5000 \
--advanced \
--minreplicashpas 2 \
--maxreplicashpas 3 \
--kafkaclustersize 3 \
--zookeeperclustersize 3 \
--couchdbclustersize 3 \
--datalayerclustersize 3 \
--cassandraclustersize 3 \
--redisServerReplicas 3 \
--cassandraUsername customCassandraSuperuser \
--masterPort 443 \
--ingressPort 443

```

- Optional: You can opt to create a `ClusterImagePolicy` to enlist the Docker registry that is used in IBM Cloud Private to the `ClusterImagePolicy` whitelist. The necessary YAML file is created automatically. Create the `ClusterImagePolicy` policy by running the following command:

```
kubectl create -f my_namespace-my_release_name-image-policy.yaml
```

- Optional: If you want to change the raw metric retention period from the default 8 days, use the `--set` option when you are running the **helm install** command in [step 15](#).

```
--set global.metric.retention.rawMaxDays=2
```

where 2 represents the number of days to retain and can be a whole number 2 - 32. Any value beyond 32 days is not recommended and can compromise Cloud App Management performance.

```
--set global.metric.summary.enabled=true
```

For more information, see [“Data retention and summarization”](#) on [page 765](#).

The following example shows a Helm installation command that sets the data retention to 15 days:

```
helm install --name ibmcloudappmgmt --values ibmcloudappmgmt.values
--set global.metric.retention.rawMaxDays=15 my_install_dir/ibm-cloud-appmgmt-
prod-1.6.0.tgz --tls
```

- Change to the `clusterAdministration` directory and run the `createSecurityClusterPrereqs.sh` script to create the `SecurityContextConstraint` for the Cloud App Management resources to use:

```
cd ~/install_dir/ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/pre-install/
clusterAdministration
./createSecurityClusterPrereqs.sh
```

- Change to the `install_dir` directory.

```
cd ~/install_dir
```

- Deploy the Cloud App Management server Helm chart using the Helm CLI.

```
helm install --name my_release_name --namespace my_namespace --values
my_release_name.values.yaml \
charts/ibm-cloud-appmgmt-prod-1.6.0.tgz --tls
```

17. After the Helm chart is successfully deployed, run the `ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/post-install-setup.sh` script to complete administrative tasks necessary to access the IBM Cloud App Management dashboard. Run the script without any parameters first to see which parameters are required and optional. For example:

```
Usage: ./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/post-install-setup.sh

--releaseName <name>          Release name, default of ${default_release}"
--namespace <name>           Namespace, default of ${namespace}"
--instanceName <name>        Name for the serviceinstance, default of ${instance_name}"

[ --advanced ]               Choose Advanced offering ( omit this parameter will chose
Base offering )"
[ --noLog ]                  Do not log to ${log_file}"
[ --tenantID <UUID> ]       The TenantID of the new serviceinstance, default is random"

"example: for Base offering"
./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/post-install-setup.sh \
--releaseName ${default_release} --instanceName ${instance_name} --namespace ${namespace}"

"example: for Advanced offering"
./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/post-install-setup.sh \
--releaseName ${default_release} --instanceName ${instance_name} --namespace ${namespace} --
advanced"
```

**Note:** After you run the `post-install-setup.sh` script, sometimes a `Connection timed out` error can be displayed. For more information about solving this issue, see [“The post-install-setup.sh script runs with the timeout error on Red Hat OpenShift environment”](#) on page 1406.

## Results

The Cloud App Management server is successfully installed. If you want to optionally verify your installation, you can run the `collectContainerLogs.sh` script, which collects the installation logs and outputs them to a diagnostic file. For more information about running this script, see [“Collecting the server logs for IBM Support”](#) on page 1417.

## What to do next

1. Start the service instance and access the Cloud App Management console. For more information, see [“Starting the Cloud App Management UI”](#) on page 176.
2. Deploy ICAM agents. For more information, see [Chapter 13, “Deploying ICAM Agents,”](#) on page 193.

## Configuring the Helm charts

After you download the Cloud App Management PPA file from the [IBM Passport Advantage](#) website, extract it, and load it into Docker, next, you must configure the Helm chart for the Cloud App Management configuration deployment that you want. Configure the Helm chart by configuring the options in the `.yaml` files. The main file that you need to edit is the `values.yaml` file.

### *Editing the configuration options in values.yaml*

To configure your Cloud App Management deployment, modify the Helm chart by editing the `values.yaml` file. Change the values of settings such as `environmentSize:` and `storageClassName:`. The `environmentSize:` setting adjusts the size of your environment. The `storageClassName:` setting changes the storage type that you want to use in your deployment.

## Before you begin

Before you edit the `values.yaml` file, you must download and extract the contents of the Cloud App Management IBM Passport Advantage Archive (PPA) file. For more information, see [“Offline: Installing IBM Cloud App Management stand-alone on IBM Cloud Private”](#) on page 148.

**Note:** Use all lowercase characters when editing the values in the `values.yaml` file.

## Procedure

1. Access the `/charts/` directory, where the PPA file is located. Open the `/my_install_dir/ibm-cloud-appmgmt-prod/values.yaml` file in a text editor of your choice.

Where `my_install_dir` is the IBM Cloud Private installation directory that you specified when you extracted the PPA file.

2. Find the setting that you need to configure. Enter the values for that setting, or replace the existing values with your own custom values. Save and close the file.
3. Issue the following command from the `my_install_dir` directory to repackage the Helm chart:

```
helm package ibm-cloud-appmgmt-prod
```

Use the following table as a reference when you are configuring the settings. It includes each option that you can edit, descriptions, the values that you can enter for some options, and examples.

Configuration setting	Value/Description/Example
<code>global.environmentSize</code>	<code>[size0_amd64   size1_amd64   -- size0_ppc64le   -- size1_ppc64le]</code> This setting determines the Kubernetes resource requests and the limits that are used by the microservices. <code>size0</code> minimises the resources. Use only for small basic tests and trials. <code>size1</code> sets the microservices to larger resource requests and limits. Use for production or larger tests and trials. For more information, see <a href="#">“Planning hardware and sizing” on page 77</a> .
<code>global.license</code>	Accept the ICAM license by replacing the empty string with "accept".
<code>global.masterIP</code>	The external IP address of the IBM Cloud Private master.
<code>global.masterPort</code>	Port address of the IBM Cloud Private master. For example: 8443.
<code>global.proxyIP</code>	The external IP address of the IBM Cloud Private proxy.
<code>global.proxyHost</code>	The full hostname address of the IBM Cloud Private proxy. For example: <code>vm1.mydomain.com</code> .
<code>global.ingress.domain</code>	The full hostname address of the IBM Cloud Private proxy. For example: <code>vm1.mydomain.com</code> .
<code>global.image.repository</code>	The Docker image repository. For example: <code>mycluster.icp:8500/default</code> .
<code>global.persistence.storageClassName</code>	The environment-wide storage class. If you are individually setting the <code>storageClassOption</code> , which is required for local storage, leave this setting empty. Example: <code>vsphere-class</code> .

Configuration setting	Value/Description/Example
global.persistence.storageClassOption.cassandradata	The storage class for the Cassandra data. If you are using an environment-wide <code>global.persistence.storageClassName</code> , such as <code>local-storage-cassandra</code> , leave this setting as default.
global.persistence.storageClassOption.zookeeperdata	Storage class for the ZooKeeper data. If you are using an environment-wide <code>global.persistence.storageClassName</code> , such as <code>local-storage-zookeeper</code> , leave this setting as default.
global.persistence.storageClassOption.kafkadata	Storage class for the Kafka data. If you are using an environment-wide <code>global.persistence.storageClassName</code> , such as <code>local-storage-kafka</code> , leave this setting as default.
global.persistence.storageClassOption.datalayerdata	Storage class for the Datalayer data. If you are using an environment-wide <code>global.persistence.storageClassName</code> , such as <code>local-storage-datalayer</code> , leave this setting as default.
global.persistence.storageClassOption.couchdbdata	Storage class for the CouchDB data. If you are using an environment-wide <code>global.persistence.storageClassName</code> , such as <code>local-storage-couchdb</code> , leave this setting as default.
global.persistence.storageSize.cassandradata	Storage size for the Cassandra data.
global.persistence.storageSize.zookeeperdata	Storage size for the ZooKeeper data.
global.persistence.storageSize.kafkadata	Storage size for the Kafka data.
global.persistence.storageSize.datalayerdata	Storage size for the Datalayer data
global.persistence.storageSize.couchdbdata	Storage size for the CouchDB data.
ibm-cem.icpbroker:adminusername	Cluster administrator user name. The default name is <code>admin</code> , which is configured when IBM Cloud Pak for Multicloud Management is installed. Replace this default name when you are using a different cluster administrator user name in your environment.

### Moving to a custom namespace

If you loaded IBM Cloud App Management on one namespace, but want to move to another namespace, modify the image scope of each Docker image.

### Procedure

1. To get the IBM Cloud App Management Docker image list, run the following command:

```
helm install --set global.license=accept --name dry-run --dry-run --debug
decompressed_ppa_file --tls |
awk -F ':' '/image:/{print $2}' | sed -e 's#["]*/##' |
sort | uniq > /tmp/icam-image-list.txt
```



Where *decompressed\_ppa\_file* is the decompressed Cloud App Management installation image file, such as the `ibm-cloud-appmgmt-prod-1.6.0.tgz` file.

2. Modify the image scope by running the following command:

```
for i in `cat /tmp/icam-image-list.txt | sed -e 's/_/-u-/g'`
do kubectl get images $i -o yaml | sed -e 's/scope: namespace/scope: global/' |
kubectl replace -f - done
```

3. Edit the `global.image.repository` setting in the `my_release_name/values.yaml` file to use the default namespace, where *my\_release\_name* is the name of your Cloud App Management Helm chart, such as `ibmcloudappmgmt`. The following example uses the `"mycluster.icp:8500/default"` default namespace:

```
global:
  environmentSize: "size0"
  imageNamePrefix: ""
  masterIP: 9.42.2.70
  masterPort: 8443
  proxyIP: 9.42.2.70
  proxyHost: "icp-213-1.rtp.raleigh.ibm.com"
  ingress:
    domain: "icp-213-1.rtp.raleigh.ibm.com"
  image:
    repository: "mycluster.icp:8500/default"
  sidecar:
    imageGroup: ""
  persistence:
    enabled: true
    storageClassName: ""
    storageClassOption:
      cassandradata: "local-storage-cassandra"
      cassandrarak: "none"
      zookeeperdata: "local-storage-zookeeper"
      kafkadata: "local-storage-kafka"
      couchdbdata: "local-storage-couchdb"
      datalayerjobs: "local-storage-datalayer"
    storageSize:
      cassandradata: "50Gi"
      cassandrarak: "50Gi"
      zookeeperdata: "1Gi"
      kafkadata: "10Gi"
      couchdbdata: "1Gi"
      datalayerjobs: "1Gi"
  ibm-cem:
    license: "accept"
```

Where `ibmcloudappmgmt` is the release name in the example.

4. Install Cloud App Management to your custom namespace, such as `kube-public`, as in the following example:

```
helm install --set global.license=accept --name ibmcloudappmgmt --values ~/
values.ibmcloudappmgmt.yaml
--tls ibm-cloud-appmgmt-prod-1.6.0.tgz --namespace kube-public
```

### Validating the Cloud App Management server deployment

Learn how to check that the Cloud App Management server is deployed successfully. Access the IBM Cloud Private console and check the release status of your server.

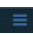
#### Procedure

Complete the following steps:

1. Open a browser window and enter the following URL to access the IBM Cloud Private console.

```
https://my_icp_console_ipaddress/console/
```

Where *my\_icp\_console\_ipaddress* is the IP address to access to the IBM Cloud Private console.

2. In the **Username** and **Password** fields, enter the IBM Cloud Private user name and password.
3. Open the  **Menu** tool in the upper left corner of the page.

4. Click **Workloads > Helm releases**.
5. In the list of Helm releases, locate the Cloud App Management server that you deployed and confirm that the server status is Deployed.
6. Optional: Select the Cloud App Management server to open the release details, scroll down to the **Pod** section, and verify that the status of all the pods is Running.  
If some pods are not running, you can check the log file for pods that have issues. Complete the following steps. From the navigation menu, click **Workloads > Deployments**; select a pod that is not running; and click the **Log** tab to display the log file for the specific pod.

## Results

You verified your Cloud App Management server deployment.

## Example

To locate the Cloud App Management server quickly, enter any unique characters of the name in the search box. The Deployed status is displayed.

NAME ^	NAMESPACE	STATUS	CHART NAME	CURRENT VERSION	AVAILABLE VERSION	UPDATED	ACTION
ibmcloudappgmt	default	Deployed	ibm-cloud-appgmt-prod	1.0.1	Up To Date	Jul 23rd 2018 08:26am	Launch

After you click the Cloud App Management server link in the **Name** field, the details of all the resources for this release are displayed in sections, such as **Pod**.

NAME	READY	STATUS	RESTARTS	AGE
ibmcloudappgmt-cassandra-0	1/1	Running	0	1d
ibmcloudappgmt-couchdb-0	1/1	Running	0	1d
ibmcloudappgmt-ibm-cem-datalayer-0	1/1	Running	0	1d
ibmcloudappgmt-kafka-0	2/2	Running	1	1d
ibmcloudappgmt-kafka-1	2/2	Running	0	1d
ibmcloudappgmt-kafka-2	2/2	Running	0	1d
ibmcloudappgmt-mongodb-0	1/1	Running	0	1d
ibmcloudappgmt-zookeeper-0	1/1	Running	0	1d
ibmcloudappgmt-zookeeper-1	1/1	Running	0	1d
ibmcloudappgmt-zookeeper-2	1/1	Running	0	1d

## Creating your service instance

After the Cloud App Management server is installed, you must create a service instance before you can access the Cloud App Management console.

## Before you begin

Ensure that you have completed the setup of persistent storage, including creating the directories on the worker nodes if local persistence is being used. For more information, see steps 6, 7 and 8 in [“Installing IBM Cloud App Management stand-alone on IBM Cloud Private”](#) on page 163.

## Procedure

To create a service instance, complete the following steps as an IBM Cloud Private cluster administrator:

1. Run the `post-install-setup.sh` script file:

```
ibm_cloud_pak/pak_extensions/post-install-setup.sh --releaseName my_release_name \  
--instanceName my_instance_name \  
--namespace my_namespace \  
[ --tenantID my_tenant_ID ]
```

Where:

- *my\_namespace* is the namespace that you selected when logging into IBM Cloud Private.
- *my\_release\_name* is the name for the Cloud App Management release that you chose during the Cloud App Management server installation. The default release name is *ibmcloudappgmt*.
- *my\_instance\_name* is the service instance name.
- *my\_tenant\_ID* is the tenantID for the Base or Advanced offering, such as 99b23e24-a751-4217-bb64-edc00b87e672. If not specified, a tenantID is randomly generated.

**Note:** Add the `--advanced` parameter to create a service instance for the IBM Cloud App Management, Advanced offering. Do not use this parameter for the IBM Cloud App Management, Base offering. For more information about the IBM Cloud App Management, Base and IBM Cloud App Management, Advanced offerings, see [“Offerings and features”](#) on page 23.

**Note:** If the stateful services do not start, ensure that you have completed all steps to setup persistent storage, including creating directories on the worker nodes if local persistence is being used. For more information, see steps 6, 7, and 8 in [“Installing IBM Cloud App Management stand-alone on IBM Cloud Private”](#) on page 163. If the issue persists, troubleshoot using the following document: [After install of Cloud App Management the stateful service pods do not start](#).

2. Run the `kubectl describe serviceinstance` command. The Cluster Service Plan External ID is displayed in the output, as in the following examples:

```
External Properties:  
Cluster Service Plan External ID: 99b23e24-a751-4217-bb64-edc00b87e672  
Cluster Service Plan External Name: base
```

or

```
External Properties:  
Cluster Service Plan External ID: 99b23e24-a751-4217-bb64-edc00b87e672  
Cluster Service Plan External Name: advanced
```

## Results

After you run the `ibm_cloud_pak/pak_extensions/post-install-setup.sh` script file, you can obtain the URL for the Cloud App Management console from the output, for example:

```
Wed Sep 11 02:46:50 EDT 2019 Done creating serviceinstance advanced  
Please access the IBM Cloud App Management dashboard at https://xxx.xxx.com/  
cemui/landing?subscriptionId=e7f18459-d3b6-40c9-b754-74e6d82b4473
```

## What to do next

After you create the service instance, you must launch it so that you are added as a user to Cloud App Management. For more information, see [“Starting the Cloud App Management UI”](#) on page 176.

## Starting the Cloud App Management UI

Log in to the Cloud App Management console from the IBM Cloud Private UI by launching your service instance. You can use the Cloud App Management console to monitor applications and services in the dashboards.

### Before you begin

**Important:** If IBM Cloud Private is configured to use a proxy external load balancer, you must ensure that the URL that you are using to log in to the Cloud App Management console includes the configured external proxy name or IP address otherwise when you try to log in to the Cloud App Management console, the page times out.

- To ensure that the user interface is not truncated, use a minimum resolution of 1280 x 1024
- For optimal performance, use the latest release of one of the following supported browsers:
  - Apple Safari
  - Google Chrome
  - Mozilla Firefox
  - Microsoft Edge

### Procedure

You can open the Cloud App Management console by entering the dashboard URL, which is obtained when creating the service instance, in a browser window. For more information, see [“Creating your service instance”](#) on page 174. You can also obtain the dashboard URL by running the following command:

```
kubectl describe serviceinstance <instance_name>  
--namespace=<namespace> | grep Dashboard | awk '{ print $3 }'
```

Where *<instance\_name>* is the service instance name and *<namespace>* is the namespace that the PPA file is loaded to.

Alternatively, complete the following steps to open the Cloud App Management console from within the IBM Cloud Private console.

1. Open a browser window and enter the following URL to access the IBM Cloud Private console:

```
https://icp_console_ipaddress/console/
```

where *icp\_console\_ipaddress* is the IP address to access to the IBM Cloud Private console.

2. In the **Username** and **Password** fields, enter the IBM Cloud Private user name and password.
3. Expand the menu in the left-hand navigation and click **Workloads > Brokered Services**.
4. Click **Launch** next to the specific service instance that you created for your Cloud App Management deployment, such as **ibmcloudapmgmt**.

### Results

After you successfully log in, the Cloud App Management Welcome page is displayed.

### What to do next

You can select the **Get Started** link to view the **Getting Started** scenarios for monitoring the health of your applications in the Cloud App Management console. For more information, see [“User interface”](#) on page 26.

You can also select **Go to my Incidents** to launch the **Incidents** page to review and act on the incidents that are assigned to you. For more information, see [“Events and incidents”](#) on page 728.

You can start deploying your agents and data collectors as described in [“Step 6: Deploy the agents and data collectors”](#) on page 70.

## Backing up and restoring

Back up your stateful services.

Kafka data is transient and does not need to be backed up.

If you are backing up metrics, the PV needs to be equal in size to your full Cassandra. Backing up metrics is not currently a recommended action.

For instructions on backing up and restoring Cassandra, see [“Back up and restore Cassandra” on page 177](#).

For instructions on backing up and restoring CouchDB, see [“Back up and restore CouchDB” on page 178](#).

### Back up and restore Cassandra

The following procedure describes how to back up and restore Cassandra.

#### Before you begin

##### Disk space

You can create a separate persistent volume for the Cassandra backup. Use the `prepare.pv.sh` script and specify a value for the **CassandraBackupDir** parameter command that is described in step [“10” on page 143](#) in the *Installing the Cloud App Management server* topic. If you are backing up metrics, the PV needs to be the same size as your full Cassandra. Backing up metrics is not currently a recommended action.

#### About this task

Back up the keyspaces individually.

#### Procedure

1. Use the following command to back up the Cassandra data (keyspaces). You need to back up only the following keyspaces: `datalayer`, `jaeger_v1_opentt`, and `subgraph`. Use the following command:

```
kubectl exec my_release_name-cassandra-number -- bash -c "/opt/ibm/backup_scripts/backup_cassandra.sh -k 'keyspace_to_backup' -f"
```

where

`my_release_name` is the name that you specified for the `--releasename` parameter during installation (`pre-install.sh`) script.

`number` is the number of the first Cassandra pod. With three pods and a replication factor of three, each pod has a copy of 100% of the data.

`keyspace_to_backup` is the individual keyspace to back up.

The best practice is to specify keyspaces individually. Backup only the following keyspaces: `datalayer`, `jaeger_v1_opentt`, and `subgraph`.

For example:

```
kubectl exec ibmcloudappmgmt-cassandra-0 -- bash -c "/opt/ibm/backup_scripts/backup_cassandra.sh -k 'datalayer' -f "
```

```
kubectl exec ibmcloudappmgmt-cassandra-0 -- bash -c "/opt/ibm/backup_scripts/backup_cassandra.sh -k 'jaeger_v1_opentt' -f"
```

```
kubectl exec ibmcloudappmgmt-cassandra-0 -- bash -c "/opt/ibm/backup_scripts/backup_cassandra.sh -k 'janusgraph' -f "
```

2. List the backup files by entering the command:

```
kubectl exec ibmcloudappmgmt-cassandra-0 -- bash -c "ls -lart /opt/ibm/cassandra/data/backup_tar"
```

The directory `/opt/ibm/cassandra/data` in the Cassandra container is mapped to the Cassandra persistent storage, example, `/k8s/data/cassandra`. The directories are the same ones that you prepared in the [“Installing IBM Cloud App Management stand-alone on IBM Cloud Private”](#) on page 163 topic.

You see a result similar to:

```
cassandra_ibmcloudappmgmt-cassandra-0_KS_datalayer_date_2019-06-11-1336-56.tar
cassandra_ibmcloudappmgmt-cassandra-0_KS_jaeger_v1_opentt_date_2019-06-11-1337-24.tar
cassandra_ibmcloudappmgmt-cassandra-0_KS_janusgraph_date_2019-06-11-1337-46.tar
```

3. Copy the files into a local directory or another persistent storage and remove the older backup files periodically. The command to copy the files from the Cassandra container to a local directory is:

```
kubectl cp ibmcloudappmgmt-cassandra-0:/opt/ibm/cassandra/data/backup_tar local_directory_name
```

The command to remove the backup files from the Cassandra container is

```
kubectl exec ibmcloudappmgmt-cassandra-0 -- bash -c "rm -f /opt/ibm/cassandra/data/backup_tar/*"
```

## Restore Cassandra

**Note:** Before you restore Cassandra, place the backup files in the `/opt/ibm/cassandra/data/backup_tar` directory in the Cassandra directory. Additionally, remove any other files that have later time stamps in the directory. The restore script uses the file that has the same keyspace name and the latest time stamp in the directory.

4. Restore Cassandra data by running the following command:

```
kubectl exec my_release_name-cassandra-number -- bash -c "/opt/ibm/backup_scripts/restore_cassandra.sh -k 'keyspace_to_restore' -f"
```

where

*my\_release\_name* is the name that you specified for the `--releasename` parameter during installation (`pre-install.sh`) script.

*number* Is the number of the first Cassandra pod. With three pods and a replication factor of three, each pod has a copy of 100% of the data.

*keyspace\_to\_restore* is the individual keyspace to restore, the best practice for Cloud App Management 2019.4.0 is to specify keyspaces individually.

Verify that the restore is successful

5. **Note:** Some events that were open before the backup can initially remain open after the restore. If the actual issue is not still occurring, the event will close automatically after a period of time.

Ensure that there are no failure messages in the restore logs, and that the Cloud APM console is displayed successfully.

## Back up and restore CouchDB

Run the procedures described to backup and restore CouchDB.

### Procedure

1. Run the following command from a command-line shell on the master node to backup CouchDB:

```
ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/backupcouch.sh -r my_release_name {-n namespace} {-o output_dir}
```

where

*my\_release\_name* is the name that you specified for the `--releasename` parameter during installation in the `pre-install.sh` script

*namespace* is the name that you specified for the **--namespace** parameter during installation in the `pre-install.sh` script

*output\_dir* specify a directory, if not specified, the output directory is `/tmp`

## 2. Run the following command to restore CouchDB

```
ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/restorecouch.sh -r my_release_name -f backup_file [-n namespace [-s y/n>]]
```

where

*my\_release\_name* is the name that you specified for the **--releasename** parameter during installation in the `pre-install.sh` script.

*backup\_file* is the name of the back up file.

*namespace* is the name that you specified for the **--namespace** parameter during installation in the `pre-install.sh` script.

Restart the CouchDB service after restoring data for the changes to take effect.

## Back up Helm release

Before you upgrade your Cloud App Management server, you might want to back up the current Helm release.

### 1. **Optional:** Back up the current Helm release of `ibm-cloud-appmgmt-prod`:

a. Use the `cloudctl` CLI to log in to the cluster as a user with the cluster administrator role, for example `cloudctl login`.

b. Fetch and save the current release contents:

```
helm get my_release_name --tls > my_release_name-date-backup.yaml
```

c. Fetch and save user-supplied values for the current release:

```
helm get values my_release_name --tls > my_release_name-date-user-values.yaml
```

d. Fetch and save the full set of values for the current release:

```
helm get values my_release_name --tls --all > my_release_name-date-all-values.yaml
```

Where *date* is set to a valid value for file names, for example, `DATE=$(date +%Y-%m-%d_%H-%M-%S)`





---

# Chapter 10. Configuring IBM Cloud App Management

This section includes configuration tasks to complete after you either install IBM Cloud App Management on its own or you install it with IBM Cloud Pak for Multicloud Management.

## Upgrading IBM Cloud App Management with IBM Cloud Pak for Multicloud Management from Eventing only to full monitoring mode

You can reconfigure your eventing only environment to support full monitoring on the hub cluster. You can complete this upgrade on your existing product version installation only, that is, you can reconfigure your V2019.4.0 IBM Cloud App Management with IBM Cloud Pak for Multicloud Management eventing only environment to support full monitoring on 2019.4.0. You cannot change to another product version during the reconfiguration.

### Before you begin

- **Important:** If you are planning on updating from eventing only to full monitoring, you must ensure to keep the storage class and storage size values the same and you must check whether your environment size can handle this update. For more information, see [“Planning hardware and sizing”](#) on page 77.
- You currently have the Cloud App Management server eventing only installation image installed. To update to full monitoring for an offline installation (an installation where you are not using Entitled Registry), you must download and install the Cloud App Management server full monitoring installation image file, complete these steps:
  1. Locate the Cloud App Management server full monitoring installation image file on IBM Passport Advantage®. Download the installation image file to the `install_dir` directory. Depending on your platform, choose one of the following images. Find the installation image by searching for it using its part number.
    - For Linux x86\_64, choose `icam_ppa_2019.4.0_prod.tar.gz` (Part number: CC4KNEN)
    - For Linux on Power, choose `icam_ppa_2019.4.0_prod_ppc64le.tar.gz` (Part number: CC4LHEN)

For more information about the IBM Cloud App Management components in the context of IBM Cloud Pak for Multicloud Management, see the IBM Cloud App Management for IBM Cloud Pak for Multicloud Management packages section in the [Passport Advantage part numbers](#) topic.

2. In the master node, change to the `install_dir` directory:

```
cd install_dir
```

3. Load the Passport Advantage Archive (PPA) file installation image file into IBM's Docker registry:

```
cloudctl catalog load-archive --archive ./installation_image_file --registry $(oc registry info)/kube-system
```

where `installation_image_file` is the compressed Cloud App Management server full monitoring installation image file that you downloaded in [step 1](#).

### Procedure

Complete the following steps to update your eventing only installation to full monitoring mode:

1. Log in to the IBM Cloud Pak for Multicloud Management console.
2. Click the menu in the upper left corner, and select **Helm Releases**.
3. Search for the `ibm-cloud-appmgmt-prod` Helm chart and select it.
4. Click **Upgrade** and expand **All parameters**.

5. Under **All parameters**, select the **Resource monitoring services enabled** checkbox and configure the following additional parameters:

Table 17. Parameters

Parameter name	Description/Commands	Example
<b>Create TLS Certs</b>	Select the checkbox.	✓
<b>Cloud Proxy FQDN</b>	<pre>oc get configmap ibmcloud-cluster- info -n kube-public - o=jsonpath='{.data.pr oxy_address}'</pre>	icp-proxy.apps.test-ocp42.os.fyre.ibm.com
<b>Cloud Proxy Client Secret</b>	It is usually <i>release_name-ingress-client</i>	ibmcloudappmgmt-ingress-client
<b>Cloud Proxy TLS Secret</b>	It is usually <i>release_name-ingress-tls</i>	ibmcloudappmgmt-ingress-tls
<b>Host Alias - Cloud Proxy</b>	<p>The IP address of the IBM Cloud Pak for Multicloud Management proxy. It is used where the DNS does not resolve the IBM Cloud Pak for Multicloud Management proxy's fully qualified domain name (FQDN). It can be determined by running:</p> <pre>kubectl get no -l proxy=true - o=jsonpath='{ \$.item s[*].status.address s[? (@.type=="InternalIP" )].address }'</pre>	10.21.17.70
<b>Product Name</b>	Change the name from Event Management for Multicloud Manager to IBM Cloud App Management for Multicloud Manager	IBM Cloud App Management for Multicloud Manager

6. Click **Upgrade**.

7. Verify that the upgrade is successful and all the pods are running:

```
kubectl get pods -l release=my_release_name
```

where *my\_release\_name* is the name that you are using for the IBM Cloud App Management release, for example: *ibmcloudappmgmt*.

8. Run the following command to upgrade the service instance and subscription:

```
kubectl exec -n my_namespace -t `kubectl get pods -l release=my_release_name -n
my_namespace -l component=cem-users | grep "Running" \
| head -n 1 | awk '{print $1}'` -- "sh" -c 'curl -d "{\"service_id\": \"941a5588-b6a2-41f2-
be9c-e7c87839cea7\", \"plan_id\": \"29a1b47b-176e-41e0-ae7e-202f489d6f01\"}" \
-H "Content-Type: application/json" -X PATCH $BROKERS_URL/icp/api/v2/service_instances/
account_id -u `kubectl get secret my_release_name-cem-brokers-cred-secret -n my_namespace
```

```
\
-o yaml | grep username | awk '{print $2}' | base64 --decode:`kubect1 get secret
my_release_name-cem-brokers-cred-secret -n my_namespace -o yaml | grep password | awk
'{print $2}' | base64 --decode`"
```

where *my\_namespace* is kube-system, *my\_release\_name* is the name that you are using for the IBM Cloud App Management release, for example: `ibmcloudappmgmt`, and *account\_id* is the `cloudctl` account ID. (Open the visual terminal (kui-shell) if the service is available to you, otherwise open your own terminal and log in using `cloudctl`. Run the `cloudctl iam accounts` command to get your account IDs.)

9. Open an interactive shell and connect to the Couchdb pod:

```
kubect1 exec -it my_release_name-couchdb-0 bash
```

10. In the Couchdb pod, for each `cloudctl` account, run the following command and substitute *account\_id* from [step 8](#).

```
curl http://localhost:5984/collabopsuser/account_id | sed 's/"offerings":\[["cem-
mcm"]/"offerings":\[["cem-apm-mcm"]/g' | curl -X PUT -d @- http://localhost:5984/
collabopsuser/account_id
```

11. Edit the `k8sdc` resource and change `collectEvents`` and `collectResources` to `on`.

```
oc edit k8sdc k8sdc-cr -n multic1uster-endpoint -o yam1
collectEvents: "on"
collectResources: "on"
```

12. In to the IBM Cloud Pak for Multicloud Management console, select the **Administration** tab from the menu bar. Access to full monitoring capabilities and the existing eventing management functions is now available.

## Results

The Cloud App Management server with IBM Cloud Pak for Multicloud Management installation is successfully updated from eventing only mode to full monitoring mode.

## What to do next

You can verify that your existing ICAM klusterlets that supported eventing only reporting are now reconfigured to support full monitoring reporting.



---

# Chapter 11. Uninstalling IBM Cloud App Management

To uninstall the Cloud App Management server, first delete the service instance and then delete the Helm deployment.

## Deleting the IBM Cloud Private service instance

Delete the IBM Cloud Private service instance before you uninstall the Cloud App Management server.

### Procedure

You can delete the service instance by running a command from the kubectl CLI or you can delete it from the IBM Cloud Private console. The steps for both methods are included.

To delete the service instance using kubectl, issue the following command from the kubectl CLI:

```
kubectl delete serviceinstance --selector release=my_release_name --namespace my_namespace
```

Where *my\_namespace* is the name of the service instance, such as the default name, `ibmcloudappmgmt` and *my\_release\_name* are the release name.

**Note:** The service instance is deleted using the kubectl command. It is safe to ignore the following error message:

```
**Error from server (BadRequest): the server rejected our request for an unknown reason**
```

If you created a service instance within the IBM Cloud Private catalog, retrieve it with the following command:

```
kubectl get serviceinstance --namespace my_namespace
```

Delete the service instance by running the following command:

```
kubectl delete serviceinstance my_instance_name --namespace my_namespace
```


Where *my\_instance\_name* is the service instance name.

Alternatively, you can delete the service instance from the IBM Cloud Private console by completing the following steps:

1. Open a browser window and enter the following URL to access the IBM Cloud Private console.

```
https://my_icp_console_ipaddress/console/
```

Where *my\_icp\_console\_ipaddress* is the IP address to access to the IBM Cloud Private console.

2. In the **Username** and **Password** fields, enter the IBM Cloud Private user name and password.
3. Open the  **Menu** tool in the upper left corner of the page.
4. Click **Workloads > Brokered Services**.
5. Select the service instance that you want to delete. From the **Actions** menu, click **Remove**.

### Results

The IBM Cloud Private service instance is deleted.

## Uninstalling IBM Cloud App Management

You can uninstall IBM Cloud App Management standalone or IBM Cloud App Management that was installed with IBM Cloud Pak for Multicloud Management using these steps.

## Before you begin

1. If you are uninstalling using the command line, you must install the Helm CLI. For instructions, see [Installing the Helm CLI \(helm\)](#).
2. You must delete the IBM Cloud Private service instance before you uninstall the Cloud App Management server. For more information, see [“Deleting the IBM Cloud Private service instance” on page 185](#).

## Procedure

Uninstall the Cloud App Management server from the console:

1. Log in to your console.
2. Click the hamburger menu in the upper left corner, and select **Helm Releases**.
3. From the list of helm releases, locate the Cloud App Management Helm release.  
If you know the name of your Cloud App Management Helm release, you can search for it by entering a keyword from the name.
4. Left-click on the **open and close list of options** menu (represented by three dots) for your Helm release and select **Delete**.

Uninstall the Cloud App Management server from the CLI:

5. Log in to your cluster. You only need to complete this step if you are uninstalling IBM Cloud App Management that was installed with IBM Cloud Pak for Multicloud Management. Ignore this step if you are uninstalling IBM Cloud App Management standalone.

```
cloudctl login -a my_cluster_URL -n kube-system --skip-ssl-validation
```

where

Where *my\_cluster\_URL* is the name that you defined for your cluster such as `https://cluster_address:443`. For future references to *masterIP*, use the value you are using for *cluster\_address*. A *cluster\_address* address example is: <https://icp-console.apps.organic-bullfrog-icp-mst.domain.com:443>.

6. Find the Helm chart that you want to uninstall from the list:

```
helm list --tls | grep ibm-cloud-appmgmt
```

7. Remove the Helm chart:

```
helm delete --purge --tls my_release_name
```

Where *my\_release\_name* is the name of your Cloud App Management Helm chart, such as `ibmcloudappgmt`.

**Note:** Some CEM datalayer-cron jobs and pods might not be deleted. This is a known issue. Manually delete any remaining jobs or pods.

8. Delete the storage classes and persistent volume storage claims (PVCs) to release the claims on the persistent data store:

```
kubectl delete storageclass --selector release=my_release_name  
kubectl delete pvc --selector release=my_release_name --namespace my_namespace  
kubectl delete pv --selector release=my_release_name --namespace my_namespace
```

where *my\_namespace* is the namespace that the IBM Passport Advantage Archive (PPA) file is loaded to.

9. Delete secrets and the cluster image policy. You only need to run the second command: **kubectl delete clusterimagepolicy....** if you are uninstalling IBM Cloud App Management that was installed with IBM Cloud Pak for Multicloud Management. Ignore this command you are uninstalling IBM Cloud App Management standalone.

```
kubectl delete secrets --selector release=my_release_name --namespace my_namespace
kubectl delete clusterimagepolicy --selector release=my_release_name --namespace
my_namespace
```

10. Optional: Back up the data on the persistent storage directories that you created on the worker nodes.
11. Optional: You can safely remove the data from the persistent storage directories that you created on the worker nodes.
12. Optional: You can remove the Cloud App Management image from IBM Cloud Private. For more information, see the [Removing an image from the console](#).

## Results

Cloud App Management server Helm chart is uninstalled. The storage configuration that was required for the installation is also deleted.





---

# Chapter 12. Upgrading the Cloud App Management server from V2019.3.0 to V2019.4.0

Upgrade your Cloud App Management server from V2019.3.0 to V2019.4.0.

## About this task

### Two different ways to upgrade

If you are upgrading a Cloud App Management server installation on IBM Cloud Pak for Multicloud Management, use the catalog upgrade that is described in step “1” on page 190.

If you are upgrading a Cloud App Management server on Red Hat OpenShift, use the command line upgrade that is described in step 2.

If you are upgrading a Cloud App Management server on IBM Cloud Private, use the command line upgrade that is described in step 2.

### Supported upgrade paths

- 2019.3.0 to 2019.4.0
- 2019.3.0.1 to 2019.4.0
- If you use an earlier version of Cloud App Management server, you must upgrade to version 2019.3.0 first.

### Permissions

A user with cluster administrator role must perform the upgrade.

### Software requirements

- oc CLI (Red Hat OpenShift only)
- cloudctl CLI
- kubectl CLI
- helm CLI
- docker CLI
- web browser with access to the cluster console

For details about installing and using CLI, see the [CLI tools guide topic](#) in the IBM Cloud Pak for Multicloud Management IBM Knowledge Center.

### Hardware requirements

- Review and verify that you meet the increased resource requirements. For more information, see [Planning hardware and sizing](#).

### Before you begin

1. Load the Cloud App Management server 2019.4.0 server installation image into the catalog:
  - a. Use the cloudctl CLI to log in to the cluster as a user with the cluster administrator role, for example, `cloudctl login`.
  - b. Use the cloudctl CLI to target the namespace where the current release of `ibm-cloud-appgmt-prod` is running:

```
cloudctl target --namespace my_namespace
```

- c. Use the docker CLI to log in to the private registry where the images will be loaded:

```
docker login $CONSOLE_CA_DOMAIN:8500
```

Alternatively, for Red Hat OpenShift clusters, use the Kubernetes cluster service for the private registry:

```
docker login docker-registry.default.svc:5000 -u $(oc whoami) -p $(oc whoami -t)
```

- d. Use the cloudctl CLI to load the Cloud App Management server 2019.4.0 PPA into the catalog. This command loads the docker images into the console private registry, and loads the helm chart into the local chart repository:

```
cloudctl catalog load-archive --archive icam_ppa_2019.4.0_prod.tar.gz
```

Alternatively, for Red Hat OpenShift clusters, use the Kubernetes cluster service for the private registry. You need to provide authentication. The provided registry value must include the correct namespace for image scoping:

```
cloudctl catalog load-archive --archive icam_ppa_2019.4.0_prod.tar.gz \
--registry docker-registry.default.svc:5000/my_namespace
```

Save the output of this command for later reference.

**Note:** This command can take 30-60+ minutes.

2. Optional: Back up the current Helm release. A backup of the Helm release can be useful for troubleshooting purposes. For more information, see [“Back up Helm release” on page 179](#).

## Procedure

### Upgrade by using the catalog

1. Complete the following steps to upgrade by using the catalog:

- a) Open the console in a web browser:

Console address command:

```
kubectl get configmap ibmcloud-cluster-info -n kube-public -
o=jsonpath='{.data.cluster_address}'
```

Console port command:

```
kubectl get configmap ibmcloud-cluster-info -n kube-public -
o=jsonpath='{.data.cluster_router_https_port}'
```

- b) Log in as a user with the cluster administrator role.
- c) Navigate to the cloud console Helm releases page:

```
https://console_address:console_port/catalog/instances
```

- d) Select the current release of `ibm-cloud-appmgmt-prod` and select **Upgrade** to open the configuration window.
- e) Select the appropriate version of the new `ibm-cloud-appmgmt-prod` chart (v1.6.0) and validate the following required configuration parameters:
  - Ensure **Reuse Values** is selected.
  - **CEM Configuration** values are unchanged.
- f) Select **Upgrade** to apply the upgrade to the project workload.

### Upgrade by using the command line

2. Complete the following steps to upgrade by using the command line:

- a) Create an upgrade directory and save or copy the PPA installation file `icam_ppa_2019.4.0_prod.tar.gz` from the [IBM Passport Advantage](#) website to the upgrade directory. Change directory to the upgrade directory:

```
mkdir -p upgrade
cd upgrade
```

- b) Extract the Helm charts from the Passport Advantage Archive (PPA) file by running the following commands from the `upgrade` directory where you saved the `icam_ppa_2019.4.0_prod.tar.gz` file:

```
tar -xvf icam_ppa_2019.4.0_prod.tar.gz charts
```

- c) In the `upgrade` directory, enter the following command:

```
helm get values my_release_name --tls > my_release_name-overrides.yaml
```

- d) You can change the default 8 days of data retention to a different value in the range 2 - 32 days. Edit the `my_release_name-overrides.yaml` file to add the following text with the exact letter casing shown and change **rawMaxDays** to a value in the range 2 - 32.

```
# Global section
global:
  metric:
    retention:
      rawMaxDays: 16
```

Before release 2019.3.0, the default setting was 32 days; starting with release 2019.3.0 the default setting is 8 days. Any value over 32 days is not supported and can degrade Cloud App Management performance. If you prefer to keep the prior release's data retention during the server upgrade and make a later decision about reducing raw data retention, set **rawMaxDays: 32**.

**Note:** If the value of **rawMaxDays** is changed, it will only affect the retention of new metrics that are stored after the upgrade. It will not affect the retention of existing data.

- e) You can also enable metrics summarization by changing `enabled: false` to `enabled: true` in the `my_release_name-overrides.yaml`

```
global:
  metric:
    summary:
      enabled: true
# Metrics summarization is enabled when set to true.
```

For more information, see [“Data retention and summarization”](#) on page 765.

- f) Review the file `my_release_name-overrides.yaml`. Ensure that the value of `global.ingress.port` is set to 443.

For example:

```
createTLScerts:false
global:
  environmentSize: size1_amd64
  image:
    repository: mycluster.icp:8500/icam
    imageNamePrefix: ""
  ingress
    clientSecret: ibmcloudappgmt-ingress-client
    domain: my-proxy.ibm.com
    port:443
    tlsSecret: ibmcloudappgmt-ingress-tls
  kafka:
    replication
```

- g) Enter the following command:

```
helm upgrade my_release_name charts/ibm-cloud-appmgmt-prod-1.6.0.tgz --values
my_release_name-overrides.yaml --tls
```

3. If you have `MetricSummaryPolicy` service enabled, restart it.

## Results

This upgrade process is asynchronous and may take 10 minutes or more. Go to the **Monitor health > Helm Releases** page to check the status of the release. Run **kubect1** on the command line or the Visual Web Terminal in the web browser to check the health of the project pods, for example:

```
kubect1 get pods --namespace my_namespace
```

The upgrade is completed when the status of the Helm release is "success" and all pods are reporting ready or completed.

---

## Chapter 13. Deploying ICAM Agents

After the Cloud App Management server is installed, you can install monitoring agents on the system where the corresponding applications that you want to monitor are located. The monitoring agents are configured to connect to the Cloud App Management server. After this connection, monitoring data for these agents can be viewed and worked on from the Cloud App Management console.

### Before you begin

Make sure that the agent requirements are met on the system where the agent will be installed. See [Cloud App Management agent requirements](#).

### About this task

**Remember:** If you have IBM Tivoli Monitoring agents, IBM Tivoli Composite Application Manager (ITCAM) agents, or Cloud APM agents that are installed to monitor your applications, do not install the Data Center Resource Agents. Instead, it is sufficient to configure your existing agents to connect to the Cloud App Management server so that you can view monitoring data on the Cloud App Management console. You can always reconnect these agents to their previous servers anytime. For more information, see [Chapter 18, “Integrating with other products,”](#) on page 675.

Deploying the Data Center Resource Agents includes the following main steps:

1. Download the agent installation images from the [IBM Passport Advantage](#) website. See [“Downloading agents and data collectors from Passport Advantage”](#) on page 194.
2. Transfer the agent installation images to an AIX or Linux system to configure the images for server connection. See [“Configuring the downloaded images”](#) on page 194.

**Important:** You must first use an AIX or Linux system to configure the agent installation images even if you want to install the agent on Windows systems only. The image configuration script is not supported on Windows systems.

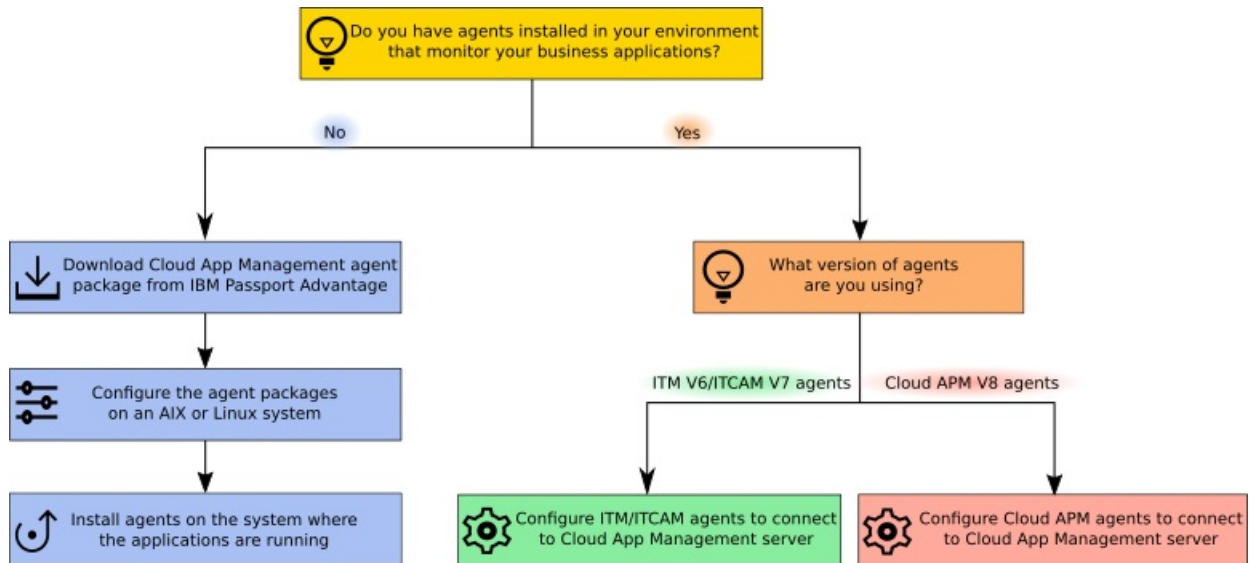
3. Install one or more agents on the system where the application that you want to monitor is located.
  - [“Installing agents on UNIX systems”](#) on page 196
  - [“Installing agents on Linux systems”](#) on page 200
  - [“Installing agents on Windows systems”](#) on page 204

**Remember:** If you install the agent as a non-root user on the AIX or Linux system, run the `UpdateAutoRun.sh` script with root user or sudo user access after a non-root installation. See [“Installing agents as a non-root user”](#) on page 209.

## Planning agent deployment

Before you can view monitoring data from the Cloud App Management UI, you must deploy monitoring agents on the system where the applications that you want to monitor are running. After that, monitoring agents can collect and send monitoring data to the Cloud App Management server for display.

Depending on whether you already have monitoring agents installed in your environment, different procedures apply. Refer to the following roadmap to find the deployment procedure that suits your environment. Click the rectangular boxes in the picture that contain the task name to get detailed information.



1. [“Downloading agents and data collectors from Passport Advantage” on page 194](#)
2. [“Configuring the downloaded images” on page 194](#)
3. [Installing the agents](#)
4. [“Integrating with IBM Tivoli Monitoring agents” on page 675](#)
5. [“Integrating with Cloud APM, Private agents” on page 688](#)

## Downloading agents and data collectors from Passport Advantage

Download the ICAM Agents and ICAM Data Collectors installation images from the [IBM Passport Advantage](#) website. Note: You must configure the ICAM Agents installation images for communication with the Cloud App Management server before you can use them to install the Data Center Resource Agents.

### Procedure

1. Log in to the [IBM Passport Advantage](#) website.
2. Identify the installation images that you want to download for the ICAM agents and ICAM data collectors.  
For part numbers and file names of each image to download, see [Table 3 on page 71](#).
3. Download the ICAM agents or the ICAM data collectors (or both) compressed installation images to an AIX or Linux system to prepare for server connection configuration.

### What to do next

- Before you install the ICAM agents, you must configure the downloaded, compressed agent installation images to communicate with the Cloud App Management server. See [“Configuring the downloaded images” on page 194](#).
- For the Data Collectors for Cloud Resources see [Chapter 15, “Deploying ICAM Data Collectors,” on page 555](#).

## Configuring the downloaded images

Before you can install ICAM Agents, you must configure the downloaded agent images for communication with the Cloud App Management server. Image configuration must be done on an AIX or Linux system. After that, use the configured agent images to install the agents on various operating systems.

## Before you begin

- If you are going to configure the Windows installation images on an AIX system, make sure the **jar**, **zip**, or **unzip** tool is available on the AIX system to process the compressed agent installation images for Windows systems.
- Check that you prepared your system to install the server in the default HTTPS mode, see [“Installing IBM Cloud App Management stand-alone on IBM Cloud Private” on page 163](#). Also, familiarise yourself with the security certificates that are available, see [“Configuring certificates for HTTPS communications” on page 152](#).

## About this task

Run the **pre\_config.sh** script on an AIX or Linux system first before you configure the agent images for all supported operating systems. The **pre\_config.sh** script is available in an agent configuration pack that can be downloaded from the Cloud App Management console.

The ICAM Agents configuration pack is populated automatically with the appropriate security features that are based on the Cloud App Management server configuration. HTTPS is enabled on the Cloud App Management server if you run the **pre-install.sh** command, or if you set either the `https` flag or the `advanced` flag during the server installation.

## Procedure

1. Download the agent configuration pack from the Cloud App Management console. The downloaded package contains agent configuration files for server connection.
  - a) Log in to the Cloud App Management console and click **Get Started**.
  - b) Click **Administration > Integrations > New integration**.
  - c) In the Standard monitoring agents section, go to the **ICAM Agents** tile and click **Configure**.
  - d) Click **Download file** to download the `ibm-cloud-icam-agents-configpack.tar` file.
2. Extract the `ibm-cloud-icam-agents-configpack.tar` file to a local system.

```
tar -xf ibm-cloud-icam-agents-configpack.tar
```

In the current directory, the `pre_config.sh` and `env.properties` files are created.

3. Run the **pre\_config.sh** script to configure all installation images.

```
./pre_config.sh -s src_images_dir -d dst_images_dir -e env.properties
```

where:

- `src_images_dir` is the local directory where the downloaded agent images are saved.
- `dst_images_dir` is the directory to output the configured agent images. If not specified, the configured agent images are saved in the `/depot` folder within the parent directory that contains the agent configuration pack. For example, if the **pre\_config.sh** is in the `/images/preconfigpack/` directory, the configuration agent images are saved in the `/images/depot/` directory.

The **pre\_config.sh** script scans the source directory to find and configure all installation images and then saves the configured images to the destination directory.

## Results

All agent installation images in the source directory are configured and the configured images are located in the destination directory. The agents or data collectors are configured to use the security certificates at run time and to communicate with the Cloud App Management server in HTTPS mode, if you choose to install the Cloud App Management server in HTTPS mode.

## What to do next

Use the configured installation images to install monitoring agents on the systems where the corresponding applications are located.

- [“Installing agents on UNIX systems” on page 196](#)
- [“Installing agents on Linux systems” on page 200](#)
- [“Installing agents on Windows systems” on page 204](#)

## Installing agents on UNIX systems

Install monitoring agents on your AIX or Solaris systems for the resources that you want to manage.

The following agents are supported on AIX systems:

- Cassandra agent
- DataPower agent
- Db2 agent
- Hadoop agent
- HTTP Server agent
- IBM Integration Bus agent
- Oracle Database agent
- SAP agent
- SAP HANA Database agent
- SAP NetWeaver Java Stack agent
- UNIX OS agent
- Sybase agent
- WebLogic agent
- WebSphere Applications agent
- WebSphere Infrastructure Manager agent
- IBM MQ(formerly WebSphere MQ) agent

The following agents are supported on Solaris systems:

- Monitoring Agent for Db2
- Monitoring Agent for HTTP Server
- Monitoring Agent for JBoss
- Monitoring Agent for MySQL
- Monitoring Agent for Oracle Database
- Monitoring Agent for SAP Applications
- Monitoring Agent for UNIX OS
- Sybase agent
- Monitoring Agent for WebSphere Applications

## Preinstallation on AIX systems

You must complete the required preinstallation tasks before you install agents on AIX systems. Some preinstallation tasks are agent-specific and other tasks apply to multiple agents.

**Important:** These requirements are in addition to the requirements identified in the Software Product Compatibility Reports.

For the current version requirements and dependencies for your agent, see [“System requirements” on page 75](#) for a link to the Software Product Compatibility Reports.

### All agents

The following preinstallation tasks are applicable to all agents:



### Non-root user installation

You must have read, write, and execute permissions for the installation directory. Otherwise, the installation is canceled. For more information about non-root user installation, see [“Installing agents as a non-root user”](#) on page 209.

### 70-character limitation for installation path

The installation directory and the path to it must be no more than 70 characters.

### AIX only: 100-character limitation for .tar file names

The default **tar** command on AIX systems cannot handle file names that are longer than 100 characters. To avoid installation issues, complete the following steps:

1. Download and install the GNU version of the **tar** command from the [AIX Toolbox for Linux Applications](#) website.
2. Make the GNU version your default **tar** command. Complete one of the following steps:

- In the *PATH* environment variable, put the following variable first:

```
export PATH=/opt/freeware/bin:$PATH
```

- Replace `/bin/tar` with symbolic link to `/opt/freeware/bin/tar`

Alternatively, upgrade to the latest version of AIX to receive the code fix for handling file names longer than 100 characters. For details, see the [TAR command Technote for AIX V6.1](#) or the [TAR command Technote for AIX V7.1](#).

### Specific agents

The following preinstallation tasks are applicable to the specified agents:

#### DataPower agent

Before the agent is installed, the prerequisite checker checks that *ulimit* is set to **unlimited** on AIX. You must run the **ulimit -d unlimited** command to ensure that the *max data segment size* system environment variable is set to **unlimited**. This agent cannot be installed on the same machine as the DataPower appliance that you want to monitor.

#### Oracle Database agent

The Oracle Java Database Connectivity (JDBC) driver that supports the monitored Oracle database versions is required. Install the Oracle JDBC driver from [Oracle Database JDBC driver downloads](#).

#### WebSphere Applications agent

Before the agent is installed, the prerequisite checker checks that *ulimit* is set to **524000** on the AIX system. You must run the **ulimit -d 524000** command to ensure that the *max data segment size* system environment variable is set to **524000**.

## Preinstallation on Solaris systems

You must complete the required preinstallation tasks before you install agents on Solaris systems. Some preinstallation tasks are agent-specific and other tasks apply to multiple agents.

**Note:** These requirements are in addition to the requirements identified in the Software Product Compatibility Reports.

For the current version requirements and dependencies for your agent, see [“System requirements”](#) on page 75 for a link to the Software Product Compatibility Reports.

### All agents

The following preinstallation tasks are applicable to all agents:

#### Non-root user installation

You must have read, write, and execute permissions for the installation directory. Otherwise, the installation is canceled. For more information about non-root user installation, see [“Installing agents as a non-root user”](#) on page 209.

### 70-character limitation for installation path

The installation directory and the path to it must be no more than 70 characters.

### 100-character limitation for .tar file names

The default **tar** command on Solaris systems cannot handle file names that are longer than 100 characters. To avoid @LongLink error issues, complete the following steps:

1. Download and install the GNU version of the **tar** command from the <http://www.gnu.org> website.
2. Make the GNU version your default **tar** command. Complete one of the following steps:

- In the *PATH* environment variable, put the following variable first:

```
export PATH=/opt/freeware/bin:$PATH
```

- Replace `/bin/tar` with symbolic link to `/opt/freeware/bin/tar`

### Setting the *CANDLEHOME* environment variable

If you used the ITM Agent Converter to install and configure an agent on the same managed system before, the *CANDLEHOME* environment variable changed to that directory where you installed the agent with the Agent Converter. Before you install and configure a native Cloud APM agent, you must set the *CANDLEHOME* environment variable to a different directory, otherwise, the native Cloud APM agent cannot start.

### Specific agents

The following preinstallation tasks are applicable to the specified agents:

#### HTTP Server agent

Install and run this agent as a root user. Use the same user ID to install and run the agent. If you install and run the agent as a non-root user, the non-root user must have the same user ID as the user who started the IBM HTTP Server. Otherwise, the agent has problems with discovering the IBM HTTP Server.

## Installing agents

You can install any combination of monitoring agents on a managed system. For example, if you install the DataPower agent to monitor DataPower Appliances in your enterprise environment, you might want to also install the UNIX OS agent. With the UNIX OS agent, you can monitor other aspects of the system, such as the overall CPU, memory, and disk.

For a list of the agents that run on AIX systems, see [“Installing agents on UNIX systems” on page 196](#).

### Before you begin

- Review the information in [“System requirements” on page 75](#) to make sure that you have the requirements for the agents you plan to install.
- Download the agents. See [“Downloading agents and data collectors from Passport Advantage” on page 194](#).
- Review the agent preinstallation tasks before you install the agents. See [“Preinstallation on AIX systems” on page 196](#).
- Configure the agent images with the connection details for the Cloud App Management server. See [“Configuring the downloaded images” on page 194](#).

**Important:** Java Runtime is installed only when the agent requires it and is not always available. Also, ksh is no longer required for agent installation, and SELinux in enforcing mode is supported.

### About this task

You can install monitoring agents as a root user or non-root user. If you do not have root privileges and you want to install a monitoring agent, you can install the agent as a non-root user, see [“Installing agents as a non-root user” on page 209](#). Also, you can install the agent as a non-root user if you are a host

administrator and you do not want to run the monitoring agent as a root user. Installation flow is the same as for a root user.

**Remember:** The default installation directory of Data Center Resource Agents is `/opt/ibm/apm/` on AIX or Linux systems. If the default directory is used by other programs, specify another directory during installation.

## Procedure

1. Transfer the configured agent images to a temporary directory on the local system where the applications that you want to monitor are located.

**Remember:** Make sure that the directory does not contain an older version of the archive file.

2. Extract the agent installation files by using the following command:

```
tar -xf ./agent_installation_files
```

where `agent_installation_files` is the agent installation file name for the current operating system.

The installation script is extracted to a directory named for the archive file and version. Agent binary and configuration-related files are extracted into subdirectories within that directory.

3. Run the installation script from the directory that is named for the archive file and version:

```
./installAPMAgents.sh
```

To install the agents in silent mode, see [“Installing agents silently” on page 208](#).

The installation program checks that the agent images were configured with parameters for connecting to the Cloud App Management server. If the agents images were not configured, the installation is stopped. You must configure the agent images and start the agent installation procedure from Step 1. See [“Configuring the downloaded images” on page 194](#).

4. Follow the prompts to complete installation.
  - a) Specify whether to install individual agents, a combination of the agents, or all of the agents.
  - b) Depending on whether you are installing or upgrading the agents, take one of the following steps:
    - If you are installing the agents, specify a different agent installation home directory or use the applicable default directory, `/opt/ibm/apm/agent`.
    - If you are upgrading the agents, after you are prompted for the agent installation home directory, enter the installation directory of the previous version of the agents.
  - c) When you are asked whether you accept the license agreement, enter 1 to accept the agreement and continue, or enter 2 to decline.

After you enter 1 (accept), a prerequisite scan of your environment starts and takes a few moments to complete. If any requirements are missing, a message directs you to a log file with the reason for the failure. A prerequisite, such as a missing library or insufficient disk space, stops the installation. You must address the failure, and start the installation script again.

5. If you installed the agents by using a non-root user ID, you must update the system startup scripts. See [“Installing agents as a non-root user” on page 209](#).
6. After installation is complete and the command line is available, you can repeat the steps in this procedure to install more monitoring agents on the managed system.

## Results

When installation completes, the selected agents are installed. The installation log file is `install_dir/logs/APPMGMT_Agents_install_date-time.log`.

## What to do next

- Configure the agent as required. If your monitoring agent requires configuration as described in [“Postinstallation tasks for the agents” on page 210](#) or if you want to review the default settings, see [Chapter 14, “Configuring the ICAM Agents,” on page 225](#).
- To start an agent, run the following command:

```
./name-agent.sh start
```

For information about the monitoring agent commands, including the *name* to use, see [“Using agent commands” on page 226](#).

- After you configure and start the agent, view the data that the agent is collecting from the Cloud App Management console. See [“Starting the Cloud App Management UI” on page 176](#).

## Installing agents on Linux systems

Install monitoring agents on your Linux systems for the resources that you want to manage.

The following agents are supported on Linux for System x systems:

- Amazon EC2 agent
- Amazon ELB agent
- Azure Compute agent
- Cassandra agent
- Cisco UCS agent
- Citrix VDI agent
- **2019.4.0.2** CouchDB agent
- DataPower agent
- Db2 agent
- DataStage agent
- Hadoop agent
- HTTP Server agent
- IBM Integration Bus agent
- JBoss agent
- Linux OS agent
- Linux KVM agent
- MariaDB agent
- MongoDB agent
- MySQL agent
- Oracle Database agent
- PostgreSQL agent
- RabbitMQ agent
- SAP agent
- SAP HANA Database agent
- SAP NetWeaver Java Stack agent
- **2019.4.0.2** Sterling Connect Direct agent
- **2019.4.0.2** Sterling File Gateway agent
- Sybase agent

- Tomcat agent
- VMware VI agent
- WebLogic agent
- WebSphere Applications agent
- WebSphere Infrastructure Manager agent
- IBM MQ(formerly WebSphere MQ) agent

## Preinstallation on Linux systems

You must complete the required preinstallation tasks before you install agents on Linux systems. Some preinstallation tasks are agent-specific and other tasks apply to multiple agents.

**Important:** These requirements are in addition to the requirements identified in the Software Product Compatibility Reports.

For the current version requirements and dependencies for your agent, see [“System requirements”](#) on page 75 for a link to the Software Product Compatibility Reports.

### All agents

The following preinstallation tasks are applicable to all agents:

#### Non-root user installation

You must have read, write, and execute permissions for the installation directory. Otherwise, the installation is canceled. For more information about non-root user installation, see [“Installing agents as a non-root user”](#) on page 209.

#### 70-character limitation for installation path

The installation directory and the path to it must be no more than 70 characters.

#### Ensure that the binary for bc is available on your system

The binary for `bc` is required to run `prereq_checker` when installing agents. But `bc` is missing on some Linux platforms, for example, SUSE Linux Enterprise 15. You must install `bc` and set it in the `PATH` environment variable. You can run `# which bc` to check. If it is available, you can see the following message:

```
# which bc
/usr/bin/bc
```

If `bc` is not found, you can see the following message:

```
# which bc
which: no bc in (/sbin:/usr/sbin:/usr/local/sbin:/root/bin:/usr/local/bin:/usr/bin:/bin)
```

### Specific operating systems

#### Red Hat Enterprise Linux (RHEL) 8

##### The `libnsl.so.1` package is needed on RHEL 8

By default, `libnsl.so.1` is not installed in Red Hat Enterprise Linux release 8.0. Without this package, no agent can be installed successfully. Have your administrator set up a yum repository for you, and then run this command:

```
yum install libnsl
```

After successful installation, you can see `/usr/lib64/libnsl.so.1`.

**Note:** The `libnsl.so.1` package is required only for agents. You do not need to do this step for data collectors.

## Bypassing the prerequisite scanner for some agents

Before the prerequisite scanner is updated for the latest supported release, for some agents, you can bypass the prerequisite scanner to have this supported release sooner. For suitable scenarios and instructions, see [“Bypassing the prerequisite scanner” on page 213](#).

**Note:** You do not need to do this step for data collectors.

### SUSE Linux Enterprise 15

#### The binary for bc is needed on SUSE Linux Enterprise 15

The binary for `bc` is required to run `prereq_checker` when installing agents. By default, `bc` is not installed on SUSE Linux Enterprise 15. You must install `bc` and set it in the `PATH` environment variable. You can run `# which bc` to check. If it is available, you can see the following message:

```
# which bc
/usr/bin/bc
```

If `bc` is not found, you can see the following message:

```
# which bc
which: no bc in (/sbin:/usr/sbin:/usr/local/sbin:/root/bin:/usr/local/bin:/usr/bin:/bin)
```

### Specific agents

The following preinstallation tasks are applicable to the specified agents:

#### DataPower agent

You must run the `ulimit -d unlimited` command to ensure that the *max data segment size* system environment variable is set to **unlimited**. This agent cannot be installed on the same machine as the DataPower Appliance that you want to monitor.

#### Microsoft SQL Server agent

To monitor a Microsoft SQL environment, the Microsoft SQL Server and Microsoft SQL ODBC driver must be installed before you install the Monitoring Agent for Microsoft SQL Server. For example, to install the ODBC driver on Red Hat Enterprise Linux, use the following command:

```
sudo yum install unixODBC
sudo yum install msodbcsql17
```

#### Oracle Database agent

The Oracle Java Database Connectivity (JDBC) driver that supports the monitored Oracle database versions is required. Install the Oracle JDBC driver from [Oracle Database JDBC driver downloads](#).

#### WebSphere Applications agent

- Before the agent is installed, the prerequisite checker checks that *ulimit* is set to **524000** on the Linux system. You must run the `ulimit -d 524000` command to ensure that the *max data segment size* system environment variable is set to **524000**.
- `libstdc++.so.5` is required when installing the WebSphere Applications agent on Linux for IBM Z. On SUSE Linux Enterprise 12 and 15 for IBM Z, `libstdc++.so.5` is not installed by default. Ask your system admin to install before deploying the WebSphere Applications agent. Otherwise, you will meet the following error:

```
KYN - WAS Monitoring Agent [version 07301410]:
```

Property	Result	Found	Expected
=====	=====	=====	=====
os.lib.libstdc++_64	FAIL	Unavailable	regex{libstdc++.so.5}

## Installing agents

You can install any combination of monitoring agents on a managed system. For example, if you install the DataPower agent to monitor DataPower Appliances in your enterprise environment, you might want to

also install the Linux OS agent. With the Linux OS agent, you can monitor other aspects of the system, such as the overall CPU, memory, and disk.

For a list of the agents that run on Linux systems, see [“Installing agents on Linux systems” on page 200](#).

### Before you begin

- Review the information in [“System requirements” on page 75](#) to make sure that you have the requirements for the agents you plan to install.
- Download the agents. See [“Downloading agents and data collectors from Passport Advantage” on page 194](#).
- Review the agent preinstallation tasks before you install the agents. See [“Preinstallation on Linux systems” on page 201](#).
- Configure the agent images with the connection details for the Cloud App Management server. See [“Configuring the downloaded images” on page 194](#).

**Important:** Java Runtime is installed only when the agent requires it and is not always available. Also, ksh is no longer required for agent installation. SELinux in enforcing mode is supported.

### About this task

You can install monitoring agents as a root user or non-root user. If you do not have root privileges and you want to install a monitoring agent, you can install the agent as a non-root user, see [“Installing agents as a non-root user” on page 209](#). Also, you can install the agent as a non-root user if you are a host administrator and you do not want to run the monitoring agent as a root user. Installation flow is the same as for a root user.

**Remember:** The default installation directory of Data Center Resource Agents is `/opt/ibm/apm/` on AIX or Linux systems. If the default directory is used by other programs, specify another directory during installation.

### Procedure

1. Transfer the configured agent images to a temporary directory on the local system where the applications that you want to monitor are located.

**Remember:** Make sure that the directory does not contain an older version of the archive file.

2. Extract the agent installation files by using the following command:

```
tar -xf ./agent_installation_files.tar
```

where `agent_installation_files` is the agent installation file name for the current operating system.

The installation script is extracted to a directory named for the archive file and version. Agent binary and configuration-related files are extracted into subdirectories within that directory.

3. Run the installation script from the directory that is named for the archive file and version:

```
./installAPMAgents.sh
```

To install the agents in silent mode, see [“Installing agents silently” on page 208](#).

The installation program checks that the agent images were configured with parameters for connecting to the Cloud App Management server. If the agents images were not configured, the installation is stopped. You must configure the agent images and start the agent installation procedure from Step 1. See [“Configuring the downloaded images” on page 194](#).

4. Follow the prompts to complete installation.
  - a) Specify whether to install individual agents, a combination of the agents, or all of the agents.
  - b) Depending on whether you are installing or upgrading the agents, take one of the following steps:

- If you are installing the agents, specify a different agent installation home directory or use the applicable default directory, `/opt/ibm/apm/agent`.
  - If you are upgrading the agents, after you are prompted for the agent installation home directory, enter the installation directory of the previous version of the agents.
- c) When you are asked whether you accept the license agreement, enter 1 to accept the agreement and continue, or enter 2 to decline.
- After you enter 1 (accept), a prerequisite scan of your environment starts and takes a few moments to complete. If any requirements are missing, a message directs you to a log file with the reason for the failure. A prerequisite, such as a missing library or insufficient disk space, stops the installation. You must address the failure, and start the installation script again.
5. If you installed the agents by using a non-root user ID, you must update the system startup scripts (see [“Installing agents as a non-root user”](#) on page 209).
  6. After installation is complete and the command line is available, you can repeat the steps in this procedure to install more monitoring agents on the managed system.

## Results

When installation completes, the selected agents are installed. The installation log file is `install_dir/logs/APPMGMT_Agents_install_date-time.log`.

## What to do next

- Configure the agent as required. If your monitoring agent requires configuration as described in [“Postinstallation tasks for the agents”](#) on page 210 or if you want to review the default settings, see [Chapter 14, “Configuring the ICAM Agents,”](#) on page 225.
- To start an agent, run the following command:

```
./name-agent.sh start
```

For information about the monitoring agent commands, including the name to use, see [“Using agent commands”](#) on page 226.

- After you configure and start the agent, view the data that the agent is collecting from the Cloud App Management console. See [“Starting the Cloud App Management UI”](#) on page 176.

## Installing agents on Windows systems

You can install some of the ICAM Agents on Windows systems.

The following monitoring agents are supported on Windows 64-bit systems. Where indicated, agents are also supported on Windows 32-bit systems.

- Amazon EC2 agent
- Amazon ELB agent
- Azure Compute agent
- Cassandra agent
- Cisco UCS agent
- Citrix VDI agent
- **2019.4.0.2** CouchDB agent
- Db2 agent
- DataStage agent
- Hadoop agent
- HTTP Server agent
- IBM Integration Bus agent



- JBoss agent
- MariaDB agent
- Microsoft Active Directory agent
- Microsoft Cluster Server agent
- Microsoft .NET agent
- Microsoft Exchange Server agent
- Microsoft Hyper-V Server agent
- Microsoft IIS agent
- Microsoft Office 365 agent
- Microsoft SharePoint Server agent
- Microsoft SQL Server agent
- NetApp Storage agent
- MySQL agent
- Oracle Database agent
- PostgreSQL agent
- RabbitMQ agent
- SAP agent
- SAP HANA Database agent
- SAP NetWeaver Java Stack agent
- Skype for Business Server agent
- **2019.4.0.2** Sterling Connect Direct agent
- **2019.4.0.2** Sterling File Gateway agent
- Sybase agent
- Tomcat agent
- VMware VI agent
- WebLogic agent
- WebSphere Applications agent
- IBM MQ(formerly WebSphere MQ) agent
- Windows OS agent\*

\* Supported on both 64-bit and 32-bit Windows systems.

## Preinstallation on Windows systems

You must complete the required preinstallation tasks before you install agents on Windows systems. Some preinstallation tasks are agent-specific and other tasks apply to multiple agents.

**Important:** These requirements are in addition to the requirements identified in the Software Product Compatibility Reports.

For the current version requirements and dependencies for your agent, see [“System requirements”](#) on page 75 for a link to the Software Product Compatibility Reports.

### All agents

The following preinstallation tasks are applicable to all agents:

## Installing from the command prompt on a local drive

Use the Windows command prompt to start the installation script. Do not use Windows PowerShell to start the installation script.

Copy the installation files to a local disk or a mapped network drive, and then start the installation script. Do not start the installation script from a network location.

Start the installation script from a new command prompt. Do not start the installation script from an existing command prompt because the command prompt might have outdated environment variables.

## Installing agents

You can install any combination of monitoring agents on a managed system. For example, if you install the IBM MQ (formerly WebSphere MQ) agent to monitor queue managers in your enterprise environment, you might want to also install the Windows OS agent. With the Windows OS agent, you can monitor other aspects of the system, such as the overall CPU, memory, and disk.

For a list of the agents that run on a Windows system, see [“Preinstallation on Windows systems” on page 205](#).

### Before you begin

- Review the information in [“System requirements” on page 75](#) to make sure that you have the requirements for the agents you plan to install.
- Download the agents. See [“Downloading agents and data collectors from Passport Advantage” on page 194](#).
- Review the agent prerequisite tasks before you install the agents. For details, see [“Preinstallation on Windows systems” on page 205](#).
- Configure the agent images with the connection details for the Cloud App Management server. See [“Configuring the downloaded images” on page 194](#).

### About this task

Ensure that you have adequate permission to run the agent installation script and agent commands. You must be logged in using one of the following user account types:

- Default Windows administrator user account
- Administrator user account
- User account, which is a member of the administrators group
- User account, which is registered as an administrator in Active Directory services

### Procedure

Complete these steps to install monitoring agents on VMs and systems where the Windows operating system is installed:

1. On your system, navigate to the directory where your configured agent images are located.
2. Extract the agent installation files to the location where you want to install the monitoring agent software.
3. Open a command prompt as administrator.
  - a) From the **Start** menu, type command in the search box.
  - b) Right-click **Command Prompt** from the list that displays and select **Run as administrator**.
4. From the command prompt, run the installation script with Administrator privileges from the directory that is named for the extracted installation file and version, for example, C:\images\APP\_MGMT\_WIN\_Agent\_Install\_2019.4.0.

```
cd extracted_image_directory
installAPMAgents.bat
```

To install the agents in silent mode, see [“Installing agents silently”](#) on page 208.

The installation program checks that the agent images were configured with parameters for connecting to the Cloud App Management server. If the agents images were not configured, the installation is stopped. You must configure the agent images and start the agent installation procedure from Step 1. See [“Configuring the downloaded images”](#) on page 194.

**Restriction:** For the WebSphere Applications agent, the Administrator privileges must be the same privileges that were used to install the WebSphere Application Server.

5. If you are installing the agents, specify a different agent installation home directory or use the applicable default directory, C:\IBM\APM.

If you are upgrading the agent, this step is skipped, and the agent installs into the previous installation directory.

**Remember:**

- The name of the installation directory cannot exceed 80 characters or contain non-ASCII, special, or double-byte characters. Directory names in the path can contain only the following characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ \_\:0123456789()~-. /.
  - When short file name creation (*8dot3Name*) is disabled, if directory names in the path contain spaces, installation is not supported.
6. When you are asked if you accept the license agreement, enter 1 to accept the agreement and continue, or enter 2 to decline.

After you enter 1 (accept), a prerequisite scan of your environment starts and takes a few moments to complete. If any requirements are missing, a message directs you to a log file with the reason for the failure. A prerequisite, such as a missing library or insufficient disk space, stops the installation. You must address the failure, and start the installation script again.

**Troubleshooting:** If the installation exits with the following message, check whether the Server service is started (**Start > Administrative Tools > Services**). If not, start the Server service and run `installAPMAgents.bat` again.

```
This script [installAPMAgents.bat] must be run as Administrator.
```

7. After installation is complete and the command line is available, you can repeat the steps in this procedure to install more monitoring agents on the managed system.

## Results

When installation completes, the selected agents are installed. The installation log file is `install_dir\logs\APPMGMT_Agents_install_date-time.log`.

## What to do next

- Configure your agents as required. If your monitoring agent requires configuration as described in [“Postinstallation tasks for the agents”](#) on page 210 or if you want to review the default settings, see [Chapter 14, “Configuring the ICAM Agents,”](#) on page 225.
- Before installing new agents, Windows installation program temporarily stops all agents currently running in the installed product location. After installation completes, the installation program restarts any stopped agents. You must manually restart any monitoring agent that is not automatically started by the installation program.
- Use one of the following methods to start the agent:
  - Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**. Right-click on an agent and click **Start**.
  - Run the following command:

```
name-agent.bat start
```

For information about the monitoring agent commands, including the name to use, see [“Using agent commands”](#) on page 226.

- After you configure and start the agent, view the data that the agent is collecting from the Cloud App Management console. See [“Starting the Cloud App Management UI”](#) on page 176.

## Installing agents silently

Installing agents silently reduces installation time. To install a monitoring agent in silent mode, you must download an agent installation image archive file from the IBM download site, preconfigure the agent images, extract the agent installation files, prepare a silent response file, and run the installation script in silent mode.

### Before you begin

1. Review the prerequisite tasks for installing the monitoring agents, and download and extract the agent installation files. For details, see [Installing agents on AIX systems](#), [Installing agents on Linux systems](#), or [Installing agents on Windows systems](#).
2. Complete the following steps to prepare a silent response file for installing agents:
  - a. Locate the silent installation file `APP_MGMT_silent_install.txt`, make a copy of this file, and open it in a text editor.
  - b. Uncomment the license agreement.
  - c. Complete one of the following steps to specify the agents that you want to install:
    - Uncomment the individual agents to be installed. For example:

```
INSTALL_AGENT=os
```

```
INSTALL_AGENT=mq
```

- Uncomment `INSTALL_AGENT=all` to install all agents.
- d. Uncomment `AGENT_HOME` and specify the directory where you want to install the agents.
- e. Save the file.

### Procedure

1. On the command line, change to the directory where you extracted the installation script and run the following command:

```
cd offering_Agent_Install_version
```

2. Run the installation command:

- **Linux** | **UNIX**

```
./installAPMAgents.sh -p path_to_silent_response_file
```

- **Windows**

```
installAPMAgents.bat -p path_to_silent_response_file
```

**Remember:** **Windows** When short file name creation (*8dot3Name*) is disabled on the Windows, if directory names in the path contain spaces, installation is not supported.

**Troubleshooting:** **Windows** The agents installation will fail on the Windows system if the prerequisite scanner cannot obtain the type of disk where the agent will be installed to. If this

occurs, you will see a fail result for the **validDestLocation** property in the installation log file. To override this issue, add `SKIP_PRECHECK=1` to the installation command:

```
installAPMAgents.bat -p path_to_silent_response_file SKIP_PRECHECK=1
```

## Results

The agents are installed.

## What to do next

Configure the agents. See the procedure and table of commands for [Linux and AIX systems](#) and for [Windows systems](#).

## Installing agents as a non-root user

If you do not have root privileges and you want to install a monitoring agent, you can install the agent as a non-root user. Also, you can install the agent as a non-root user if you are a host administrator and you do not want to run the monitoring agent as a root user. Installation flow is the same as for a root user. After a non-root installation, run the **UpdateAutoRun.sh** script with root user or sudo user access.

### Before you begin

To uniquely identify the computer system, the Linux OS agent must identify the computer system board Universal Unique Identifier (UUID), manufacturer, model and serial number.

To obtain the computer system information, complete the following steps:

1. Ensure that the **/usr/bin/hal-get-property** command is installed on the computer system and that the hald process (HAL daemon) is running.
2. If the **/usr/bin/hal-get-property** command is not installed on the computer system, then confirm that the `/sys/class/dmi/id/product_uuid` file exists and contains the computer system UUID.

**Note:** The Linux OS Agent does not support monitoring of Docker when running as non-root.

**Remember:** The Linux OS agent retrieves this information periodically so the commands or files in the previous steps must remain in place even after installation.

### Procedure

1. Install your monitoring agents on Linux or AIX systems, as described in [“Installing agents on Linux systems”](#) on page 200 and [“Installing agents on UNIX systems”](#) on page 196.
2. Optional: If you installed your agent as a selected user and want to configure the agent as a different user, run the `install_dir/bin/secure.sh` script.

For example:

```
install_dir/bin/secure.sh -g mqadmin1
```

For more information about the **./secure.sh** script, see [“Configuring agents as a non-root user”](#) on page 231 and [Securing the agent installation files](#).

3. Optional: Configure your monitoring agents on Linux or AIX as necessary, see [Chapter 14, “Configuring the ICAM Agents,”](#) on page 225.
4. To update the system startup scripts, run the following script with root user or sudo user access:

```
install_dir/bin/UpdateAutoRun.sh
```

## What to do next

If you installed your agent as a non-root user and you want to configure the agent as the same user, no special action is required. If you installed your agent as a selected user and want to configure the agent as a different user, see [“Configuring agents as a non-root user”](#) on page 231.

If you installed and configured your agent as a non-root user and you want to start the agent as the same user, no special action is required. If you installed and configured your agent as a selected user and want to start the agent as a different user, see [“Starting agents as a non-root user”](#) on page 230.

Use the same user ID for agent installation and upgrades.

If you run the **UpdateAutoRun.sh** script as root user, the agent is configured to automatically start after operating system restart. If you do not want this agent behavior, you can disable the automatic agent start. For more information, see [“Disabling automatic agent start on AIX and Linux systems”](#) on page 232.

## Postinstallation tasks for the agents

After installation, some agents are configured and started automatically, while some agents must be configured and started manually. Multiple instance agents require creating a first instance and starting manually.

To determine the Cloud App Management agent configuration, startup, and instance characteristics, see [Table 18](#) on page 210.

For information about how to configure an agent, see [Chapter 14, “Configuring the ICAM Agents,”](#) on page 225.

*Table 18. Agent postinstallation checklist*

Agent	Configured and started automatically	Configured manually and started automatically	Configured and started manually	Multiple instance (started manually)
Amazon EC2 agent	—	—	—	✓
Amazon ELB agent	—	—	—	✓
Azure Compute agent	—	—	—	✓
Cassandra agent	—	—	✓	✓
Cisco UCS agent	—	✓	✓	✓
Citrix VDI agent	—	—	—	✓
CouchDB agent	—	—	✓	✓
DataPower agent	—	—	—	✓
DataStage agent	—	—	✓	✓
Db2 agent	—	—	—	✓
IBM Integration Bus agent	—	—	—	✓
JBoss agent	—	—	—	✓
Linux OS agent	✓	—	—	—
Linux KVM agent	—	—	✓	✓
Microsoft Active Directory agent	✓	✓	—	—
Microsoft Cluster Server agent	—	—	✓	—

Table 18. Agent postinstallation checklist (continued)

Agent	Configured and started automatically	Configured manually and started automatically	Configured and started manually	Multiple instance (started manually)
Microsoft Exchange Server agent	—	✓	✓	—
Microsoft Office 365 agent	—	—	✓	✓
Microsoft SharePoint Server agent	✓	✓	—	—
Microsoft Hyper-V Server agent	✓	—	—	—
Microsoft IIS agent	✓	—	—	—
Microsoft SQL Server agent	—	—	✓	✓
NetApp Storage agent	—	—	✓	✓
MySQL agent	—	—	✓	✓
PostgreSQL agent	—	—	✓	✓
Oracle Database agent	—	—	—	✓
RabbitMQ agent	—	—	✓	✓
SAP agent	—	—	✓	✓
SAP HANA Database agent	—	—	✓	✓
SAP NetWeaver Java Stack agent	—	—	✓	✓
Skype for Business Server agent	—	✓	✓	✓
Sterling Connect Direct agent	—	—	✓	✓
Sterling File Gateway agent	—	—	✓	✓
Sybase agent	—	—	✓	✓
Tomcat agent	—	—	✓	✓
UNIX OS agent	✓	—	—	—
VMware VI agent	—	—	✓	✓
WebLogic agent	—	—	—	✓

Table 18. Agent postinstallation checklist (continued)

Agent	Configured and started automatically	Configured manually and started automatically	Configured and started manually	Multiple instance (started manually)
WebSphere Applications agent	—	✓ The agent is started automatically but the data collector must be configured before data is reported.	—	—
WebSphere Infrastructure Manager agent	—	—	✓	✓
IBM MQ(formerly WebSphere MQ) agent	—	—	—	✓
Windows OS agent	✓	—	—	—

## Securing the agent installation files

After you install monitoring agents as a non-root user on Linux or AIX systems, you can run the `secure.sh` script to secure the agent installation by removing world write permissions and setting correct file ownership.

### Before you begin

- You must have read, write, and execute permissions for the installation directory.
- Installation of the monitoring agents and any agent configuration must be completed on the system and the agents must be successfully started.
- If you are running agents as different user accounts, they must be members of the same group. (See the [-g option](#).)

### About this task

Complete this step to lock down the file permissions in your installation. Options are available to require no root password, to specify a group name, and to view help for the command.

### Procedure

- Run the following command from the `install_dir/bin` directory.

```
secure.sh [-g common_group] [-n] [-h]
```

- In the simplest mode, run the `./secure.sh` script, which removes world write permissions, and sets the current user and user's group as the file owners. If the script is run by a non-root user, the user is prompted for the root password.
- If a non-root user runs the `./secure.sh` script with the `-n` option, this user is not prompted for a root password. In this case, changing file permissions and changing ownership are done by using this user's privileges. If the installation directory contains files that are owned by different users and the current user has no privileges to modify permissions and ownership of other user's files, this mode can fail.



- If you want to set a certain group as the group owner, the owner must provide the `-g` option with a valid group name as an argument to that option. (See [Example](#).)  
Run `secure.sh -g common_group`.  
The command changes ownership of the files and directories recursively.

If the `common_group` group is not the user's primary group, you can set the `common_group` group to be the group owner of new files created in a directory, by running the following command:

```
chmod g+s install_dir/sub_dir
```

where, `sub_dir` is any sub-directory, for example, `/opt/ibm/apm/agent`.

- Run the `./secure.sh` script with the `-h` option to get help information for the script.

## Results

The installation directory allows access to only the user who ran the script or to only the users in the specified group.

## Example

If user Alice is a member of the system group that is named "apmgroup", she can use the group to set file group ownership with the following command:

```
./secure.sh -g apmgroup
```

After the script is run, the group is set as "apmgroup" for all files in `install_dir` for the group.

## What to do next

Running the `./secure.sh` script should result in the following permissions being set for the agents.

```
lwx lwx ---
```

After you run the script, check the permissions for the agent files. For example, for IBM MQ (formerly WebSphere MQ) agent, check the files in the `install_dir/arch/mq/lib` directory. If the permissions for these files are not set correctly, update the permissions manually. For example, for the IBM MQ (formerly WebSphere MQ) agent:

1. Set the permissions by running the following command:

```
chmod g+rx install_dir/bin/mq-agent.sh
```

2. Set the user and group by running the following command:

```
chown newuser:newgroup install_dir/bin/mq-agent.sh
```

## Bypassing the prerequisite scanner

When you install monitoring agents, a prerequisite scan of your environment starts and takes a few moments to complete. If any requirements are missing, a message directs you to a log file with the reason for the failure. In some installation scenarios, you might want to either ignore warning messages or completely bypass the prerequisite check.

### About this task

There are two levels of failure messages, WARN and FAIL, and there are two levels of bypassing:

- Setting the **IGNORE\_PRECHECK\_WARNING** variable causes the installer to ignore the warning (WARN) messages.
- Setting the **SKIP\_PRECHECK** variable causes the installer to ignore all failure messages.

If your agent installation failed and you received a warning (WARN) from the prerequisite checker, review the warning. If you want to continue with the installation, set **IGNORE\_PRECHECK\_WARNING** and install again.

In an environment where you have virtual machine images that serve as templates, the prerequisite scan that is undertaken before installation begins can be done on only the first template image. If a VM image passes the scan, the other VMs created from that image will also pass. You can save time by bypassing the prerequisite check for other VMs that were created from the same image. Set **SKIP\_PRECHECK** variable and install again.

The **SKIP\_PRECHECK** setting is also appropriate for the scenario where you have a new operating system that IBM Support or the Software Product Compatibility Reports indicate that it is supported but the prerequisite checker has not yet been updated. Be sure to first try to install the agent, check the log, and make sure that this new OS is the only item failing – and the only item that you are bypassing – because **SKIP\_PRECHECK** causes the installer to bypass every item in the prerequisite checklist.

After downloading and extracting the installation files, complete this procedure to ignore the warning messages or to bypass the prerequisite scan.

### Procedure

On the system where you plan to install monitoring agents, enter one of the following commands:

- Ignore the warning (WARN) messages during the prerequisite check:
  - **Linux** | **UNIX** `export IGNORE_PRECHECK_WARNING=1`
  - **Windows** `set IGNORE_PRECHECK_WARNING=1`
- Bypass the prerequisite scan:
  - **Linux** | **UNIX** `export SKIP_PRECHECK=1`
  - **Windows** `set SKIP_PRECHECK=1`

### What to do next

To restore the default setting the next time you want to install the agent with the prerequisite scanner, turn off the **IGNORE\_PRECHECK\_WARNING** or **SKIP\_PRECHECK** variable:

- **Linux** | **UNIX** `unset IGNORE_PRECHECK_WARNING`
- **Windows** `set IGNORE_PRECHECK_WARNING=`

or

- **Linux** | **UNIX** `unset SKIP_PRECHECK`
- **Windows** `set SKIP_PRECHECK=`

## Uninstalling your agents

Uninstall a single agent or all the agents from a managed system.

### Before you begin

For multi-instance agents, you must remove all agent instances before you uninstall the agent. Otherwise, agent entries are not cleared from the registry. To remove instances, run the following command:

- **Windows** `name-agent.bat remove instance_name`
- **Linux** | **UNIX** `./name-agent.sh remove instance_name`

where, *name* is the name of the agent and *instance\_name* is the instance name. For more information, see [“Using agent commands” on page 226](#). For a list of multiple-instance agents, see [Table 18 on page 210](#).

For the following agents, an agent-specific task must be completed before you complete the uninstallation procedure:

- For the WebSphere Applications agent, you must unconfigure the data collector for all monitored server instances before you uninstall the agent. Follow the instructions in [“WebSphere Applications agent: Unconfiguring the data collector”](#) on page 216.

For the WebSphere Applications agent, make sure that the user ID, which is used to uninstall the agent, has full read and write permissions to the `logs` and `runtime` directories and all their contained subdirectories and files within the data collector home directory. The data collector home directory is as follows:

- **Windows** `install_dir\dchome\7.3.0.14.09`
- **Linux** | **UNIX** `install_dir/yndchome/7.3.0.14.09`

### About this task

The Oracle Database agent on Windows systems can be uninstalled only by using the command prompt.

### Procedure

1. On the VM or system where the monitoring agent (or agents) is installed, start a command line and change to the binary directory:

- **Linux** | **UNIX** `install_dir/bin`
- **Windows** `install_dir\BIN`

where `install_dir` the installation directory of the monitoring agent or agents.

2. Use one of the following commands to uninstall one or more monitoring agents.

- To uninstall a specific monitoring agent, enter the agent script name and the uninstall option where `name` is the agent script name:

```
– Linux | UNIX  
  
./name-agent.sh uninstall
```

```
– Windows  
  
name-agent.bat uninstall
```

For a list of the agent script names, see [“Using agent commands”](#) on page 226.

- To uninstall all the monitoring agents from the managed system with a confirmation prompt, enter the script name and uninstall all option:

```
– Linux | UNIX  
  
./smai-agent.sh uninstall_all
```

```
– Windows  
  
smai-agent.bat uninstall_all
```

- **Linux** | **UNIX** On Linux and AIX systems, to force the uninstallation of all the monitoring agents without a prompt for confirmation, enter the script name and the force uninstall all option:

```
./smai-agent.sh uninstall_all force
```

### Results

The monitoring agents are uninstalled from the system or VM.

## WebSphere Applications agent: Unconfiguring the data collector

If you uninstall the WebSphere Applications agent before you unconfigure the data collector, the agent uninstallation fails. You can remove the data collector from an application server instance manually or by using the interactive utility or the silent unconfiguration process.

For instances that are monitored with PMI resource monitoring, unconfiguration is not available. Monitoring of PMI data continues while the server is available.

### Unconfiguring the data collector interactively

If you no longer want the data collector to monitor one or more application server instances, you can unconfigure the data collector for them.

#### Before you begin

Use the user ID for configuring the data collector to unconfigure the data collector, which is also the user ID for installing the application server. Verify that this user ID has read and write permissions to the data collector home directory and all its sub-directories. The data collector home directory is as follows, where *install\_dir* is the WebSphere Applications agent installation directory.

**Note:** The exact version of the data collector may be updated or change. For example: 7.3.0.10, 7.3.0.14.09

- **Windows** `install_dir\dchome\7.3.0.14.09`
- **Linux** | **UNIX** `install_dir/yndchome/7.3.0.14.09`

#### About this task

The unconfiguration utility (`unconfig.sh` or `unconfig.bat`) is a menu driven command-line utility for unconfiguring the data collector.

#### Procedure

1. Log in to the system as the user ID that is used to configure the data collector.
2. Navigate to the following bin directory:
  - **Windows** `agent_install_dir\dchome\7.3.0.14.09\bin`
  - **Linux** | **UNIX** `agent_install_dir/yndchome/7.3.0.14.09/bin`
3. Optional: Set the location of the Java home directory before you start the utility.  
For example:

```
Linux | UNIX export JAVA_HOME=/opt/IBM/AppServer80/java
```

```
Windows set JAVA_HOME=C:\Progra~1\IBM\WebSphere\AppServer80\java
```

4. Start the unconfiguration utility by issue the following command:

```
Linux | UNIX ./unconfig.sh
```

```
Windows unconfig.bat
```

5. The utility searches for all server instances that are monitored by the data collector. Enter the number that corresponds to the application server instance to unconfigure for data collection or enter an asterisk (\*) to unconfigure data collection for all application server instances. To specify a subset of servers, enter the numbers, separated by commas, that represent the servers. For example: 1,2,3.

#### Remember:

- For a stand-alone environment, application server instances must be running during the configuration. (A WebSphere Application Server Liberty instance does not need to be running).

- For a Network Deployment environment, the Node Agent and Deployment Manager must be running.
6. The utility prompts you to specify whether you want to create a backup of your current WebSphere Application Server configuration. Enter 1 to create a backup of the current configuration. Otherwise, enter 2 and skip to step 8.
  7. The utility prompts you to specify the directory in which to store the backup of the configuration. Specify a directory in which to store the backup of the configuration or accept the default directory. The utility displays the name of the WebSphere home directory and the WebSphere profile for which a backup is created.
  8. The utility indicates whether WebSphere Global Security is enabled for the WebSphere Application profile that you specified. If global security is not enabled, skip to step 10.
  9. The utility prompts you to specify whether to retrieve security settings from a client properties file. Enter 1 to allow the utility to retrieve the user name and password from the appropriate client properties file and skip to step “10” on page 217. Otherwise, enter 2 to enter the user name and password.

The data collector communicates with the WebSphere Administrative Services using the RMI or the SOAP protocol. If global security is enabled for a profile, you must specify the user ID and password of a user who is authorized to log in to the IBM WebSphere Application Server administrative console for the profile. Alternatively, you can encrypt the user name and password and store them in client properties files before configuring the data collector. You must use the `sas.client.props` file for an RMI connection, or the `soap.client.props` file for an SOAP connection.

If you selected the option to back up the current WebSphere configuration, the utility starts backing up the configuration.

10. The utility unconfigures the data collector for the specified application server instances. A status message is displayed to indicate that the data collector was successfully unconfigured.
11. After the data collector unconfiguration completes, restart the application server instances. The data collector configuration takes effect when the application server instances are restarted. PMI resource monitoring for the server instance is still available.
12. Optional: If you want to use resource monitoring for a server instance after unconfiguring the data collector, restart the monitoring agent by running the following commands:

- **Windows**

```
cd install_dir\bin
was-agent.bat stop
was-agent.bat start
```

- **Linux | UNIX**

```
cd install_dir/bin
./was-agent.sh stop
./was-agent.sh start
```

## Results

The data collector is unconfigured for the specified application server instances.

### Unconfiguring the data collector in silent mode

You can unconfigure the data collector using the unconfiguration utility in silent mode.

### Before you begin

Use the user ID for configuring the data collector to unconfigure the data collector, which is also the user ID for installing the application server. Verify that this user ID has read and write permissions to the data collector home directory and all its sub-directories. The data collector home directory is as follows, where `install_dir` is the WebSphere Applications agent installation directory.

**Note:** The exact version of the data collector may be updated or change. For example: 7.3.0.10, 7.3.0.14.09

- **Windows** `install_dir\dchome\7.3.0.14.09`
- **Linux** | **UNIX** `install_dir/yndchome/7.3.0.14.09`

### About this task

When you unconfigure the data collector in silent mode, you first specify configuration options in a properties file. A sample properties file, `sample_silent_unconfig.txt`, is packaged with the unconfiguration utility. The file is available in `bin` directory within data collector home directory.

### Procedure

1. Log in to the system with the user ID that is used to configure the data collector.
2. Specify the configuration options in the properties `.txt` file.

The following properties are available for unconfiguring the data collector in silent mode:

#### WebSphere Application Server connecting settings

**was.wsadmin.connection.host**

Specifies the name of the host to which the wsadmin tool is connecting.

#### WebSphere Application Server global security settings

**was.wsadmin.username**

Specifies the user ID of a user who is authorized to log on to the WebSphere Application Server administrative console. This user must have the agent role on the application server.

**was.wsadmin.password**

Specifies the password that corresponds to the user specified in the `was.wsadmin.username` property.

#### WebSphere Application Server settings

**was.appserver.profile.name**

Specifies the name of the application server profile you want to unconfigure.

**was.appserver.home**

Specifies the WebSphere Application Server home directory.

**was.appserver.cell.name**

Specifies the WebSphere Application Server cell name.

**was.appserver.node.name**

Specifies the WebSphere Application Server node name.

#### Backup of the WebSphere Application Server configuration

**was.backup.configuration**

Specifies whether to back up the current configuration of the WebSphere Application Server data collector configuration before unconfiguring the data collector. Valid values are `True` and `False`.

**was.backup.configuration.dir**

Specifies the location of the backup directory.

#### WebSphere Application Server runtime instance settings

**was.appserver.server.name**

Specifies an application server instance within the application server profile for which you want to unconfigure the data collector.

**Tip:** The silent response file can have multiple instances of this property.

3. Navigate to the following directory:

- **Windows** `install_dir\dchome\7.3.0.14.09\bin`

- **Linux** | **UNIX** `install_dir/yndchome/7.3.0.14.09/bin`

4. Run the following command:

- **Windows**

```
unconfig.bat -silent path_to_silent_file
```

- **Linux** | **UNIX**

```
unconfig.sh -silent path_to_silent_file
```

5. After the data collector unconfiguration completes, restart the application server instances.

The data collector configuration takes effect when the application server instances are restarted. PMI resource monitoring for the server instance is still available.

6. Optional: If you want to use resource monitoring for a server instance after unconfiguring the data collector, restart the monitoring agent by running the following commands:

- **Windows**

```
cd install_dir\bin
was-agent.bat stop
was-agent.bat start
```

- **Linux** | **UNIX**

```
cd install_dir/bin
./was-agent.sh stop
./was-agent.sh start
```

### Manually unconfigure the data collector

After you manually configure the data collector for the WebSphere Applications agent, to remove data collection within the configured application server, you must manually unconfigure the data collector.

#### About this task

The following procedure applies only after you manually configure the data collector following the instructions in [“Manually configure the data collector if the configuration utilities fail”](#) on page 542. If you used the configuration utilities to configure the data collector, you must also use the unconfiguration utility to unconfigure the data collector. For instructions, see [“Unconfiguring the data collector interactively”](#) on page 216 or [“Unconfiguring the data collector in silent mode”](#) on page 217.

#### Procedure

- To manually unconfigure the data collector for the WebSphere application server, see [“Manually unconfiguring the data collector for WebSphere Application Server traditional”](#) on page 219.
- To manually unconfigure the data collector for the Liberty server, see [“Manually unconfiguring the data collector for WebSphere Application Server Liberty”](#) on page 220.

#### *Manually unconfiguring the data collector for WebSphere Application Server traditional*

#### Procedure

1. Log in to the WebSphere Administrative Console as the administrator.
2. In the navigation pane, click **Servers**, expand **Server Type** and select **WebSphere application servers**.
3. Click the name of the application server.
4. Under the **Server Infrastructure** section in the Configuration tab, expand **Java Virtual Machine** and click **Process Definition**.
5. Under the **Additional Properties** section, click **Java Virtual Machine**.

6. In the **Generic JVM arguments** field, remove the following entries from the content.

```
-agentlib:am_ibm_16=${WAS_SERVER_NAME} -Xbootclasspath/p:${ITCAMDCHOME}/  
toolkit/lib/bcm-bootstrap.jar -Djava.security.policy=${ITCAMDCHOME}/itcamdc/  
etc/datacollector.policy -verbosegc
```

7. Click **Apply** and click **Save**. In the Save to Master Configuration dialog box, complete the following steps:

- If you are under a Network Deployment environment, ensure that **Synchronize changes with Nodes** is selected and then click **Save**.
- If you are not under a Network Deployment environment, click **Save**.

8. In the navigation pane, click **Servers**, expand **Server Types**, click **WebSphere application servers** and then click the server name.

9. In the Configuration tab, go to **Server Infrastructure > Java and Process Management > Process Definition > Environment Entries**.

10. Depending on the operating system, the hardware platform, and the application server JVM, remove the following environment entry.

- **AIX** LIBPATH
- **Linux** LD\_LIBRARY\_PATH
- **Windows** PATH

11. In the navigation pane, click **Environment > WebSphere Variables**.

12. Remove the *ITCAMDCHOME* variable if it exists.

13. Click **Apply** and click **Save**. In the Save to Master Configuration dialog box, complete the following steps:

- If you are under a Network Deployment environment, ensure that **Synchronize changes with Nodes** is selected and then click **Save**.
- If you are not under a Network Deployment environment, click **Save**.

14. Restart the application server instance.

15. Go to the *runtime* directory in the agent installation directory and remove the *profile\_name.cell\_name.node\_name.server\_name.manual.input.properties* file.

- **Linux** **UNIX** *install\_dir/yndchome/7.3.0.14.09/runtime/  
profile\_name.cell\_name.node\_name.server\_name.manual.input.  
properties*
- **Windows** *install\_dir\dchome\7.3.0.14.09\runtime  
\profile\_name.cell\_name.node\_name.server\_name.manual.input.  
properties*

**Note:** The exact version of the data collector may be updated or change. For example: 7.3.0.10, 7.3.0.14.09

### *Manually unconfiguring the data collector for WebSphere Application Server Liberty*

#### **Procedure**

1. Navigate to the liberty server directory and open the *jvm.options* file in the *server\_name* directory within the Liberty server installation directory. For example, */opt/ibm/wlp/usr/servers/defaultServer*.
2. Remove the following parameters from the *jvm.options* file.

```
-agentlib:am_ibm_16=server_name  
-Xbootclasspath/p:dc_home/toolkit/lib/bcm-bootstrap.jar  
-Djava.security.policy=dc_home/itcamdc/etc/datacollector.policy  
-verbosegc
```



where, *server\_name* is the name of the Liberty server; *dc\_home* is the data collector home directory.

3. Open the `server.xml` file and remove the following lines:

```
<feature>webProfile-7.0</feature>
<feature>monitor-1.0</feature>
<feature>usr:itcam-7.3.0.14.09</feature>
```

4. Open the `server.env` file and remove the following entry value from the environment entry per the operating system:

Table 19. Environment entry		
Platform	Environment entry name	Environment entry value
AIX R6.1 (64-bit JVM)	LIBPATH	/lib:dc_home/ toolkit/lib/aix536
AIX R7.1 (64 bit JVM)	LIBPATH	/lib:dc_home/ toolkit/lib/aix536
Linux x86_64 R2.6 (64-bit JVM)	LD_LIBRARY_PATH	/lib:dc_home/ toolkit/lib/lx8266
Linux Intel R2.6 (32-bit JVM)	LD_LIBRARY_PATH	/lib:dc_home/ toolkit/lib/li6263
Windows (32-bit JVM)	PATH	/lib;dc_home/ toolkit/lib/win32
Windows (64-bit JVM)	PATH	/lib;dc_home/ toolkit/lib/win64

5. Restart the Liberty server.

6. Go to the runtime directory in the WebSphere Applications agent installation directory and remove the `cell_name.node_name.server_name.manual.input.properties` file.

- **Linux** `install_dir/yndchome/7.3.0.14.09/runtime/cell_name.node_name.server_name.manual.input.properties`
- **Windows** `install_dir\dchome\7.3.0.14.09\runtime\cell_name.node_name.server_name.manual.input.properties`

**Note:** The exact version of the data collector may be updated or change. For example: 7.3.0.10, 7.3.0.14.09

### Manually removing data collector configuration from an application server instance

To manually remove the data collector configuration from an application server instance, you must be able to connect to the application server by using the `wsadmin` tool. This is possible only if you are using WebSphere Application Server Network Deployment and the Deployment Manager is running. If the WebSphere application server cannot start, you must restore the WebSphere application server from the backup taken when you run the configuration utility.

#### About this task

You can manually remove the data collector configuration from an application server instance, if any of the following conditions apply:

- In a non-Network Deployment environment, you manually added the data collector configuration to the application server instance and you want to unconfigure data collection. The application server instance must be running.
- In a Network Deployment environment, you manually added the data collector configuration to the application server instance and you want to unconfigure data collection. The Node Agent and Deployment Manager on the application server must be running.

- In a Network Deployment environment, you configured the application server instance for data collection manually and the application server fails to start. The Node Agent and Deployment Manager on the application server must be running.

If you configured a stand-alone application server instance for data collection either manually or with the configuration or migration utility and the application server fails to start, you must restore your WebSphere Application Server configuration with your backup configuration. For more information, see [“Restoring the application server configuration from a backup” on page 546](#).

#### Remember:

- You must make manual changes to the WebSphere application server configuration for data collectors as the WebSphere administrative user.
- Making manual changes to the WebSphere application server for data collection must be performed by an experienced WebSphere administrator only. Any error in the manual configuration change can result in the application server not starting.
- If you manually configure the data collector to monitor application server instances, you cannot use the unconfiguration utility to unconfigure the data collector.

#### Procedure

To manually remove the data collector configuration, complete the following procedure:

1. Log in to the WebSphere Administration Server Console.
2. Click **Servers**.
3. Expand **Server Type** and select **WebSphere application servers**.
4. Click the name of the server.
5. In the Configuration tab, go to **Server Infrastructure > Java and Process Management > Process Definition > Java Virtual Machine > Additional Properties: Custom Properties**.
6. Remove any of the following JVM Custom Properties, if they are present:
  - `am.home`
  - `ITCAM.DC.ENABLED`
  - `TEMAGCCollector.gclog.path`
  - `com.ibm.tivoli.itcam.toolkit.ai.runtimebuilder.enable.rebuild`
  - `com.ibm.tivoli.jiti.injector.ProbeInjectorManagerChain.primaryInjectorFile`
7. Identify the JVM arguments that were added for the data collector.
  - a) In the navigation pane, click **Environment > WebSphere Variables**.
  - b) If you manually configured the application server for data collection, locate the JVM arguments you added manually.  
If you configured the application server for data collection with the configuration utilities, compare the values of the **AM\_OLD\_ARGS** and **AM\_CONFIG\_JVM\_ARGS** arguments to determine which arguments were added by the configuration utility.
8. Click **Server > Application Server** and select the appropriate server name.
9. In the Configuration tab, go to **Server Infrastructure > Java and Process Management > Process Definition > Java Virtual Machine**.
10. In **Generic JVM Arguments** field, remove the JVM arguments that you identified in Step 7 for the data collector.
11. Click **Apply** or **OK**.
12. In the **Messages** dialog box, click **Save**.
13. In the **Save to Master Configuration** dialog box, complete one of the following steps:
  - If you are under a Network Deployment environment, make sure that the **Synchronize changes with Nodes** check box is selected, and then click **Save**.
  - If you are not under a Network Deployment environment, click **Save**.

14. Remove environment entries that were added for the data collector.
  - a) In the Configuration tab, go to **Server Infrastructure > Java and Process Management > Process Definition > Environment Entries**.
  - b) Depending on the operating system, delete the following environment entry:
    - **AIX LIBPATH**
    - **Linux LD\_LIBRARY\_PATH**
    - **Windows PATH**
  - c) Remove the **NLSPATH** environment entry.
15. Click **Apply** or **OK**.
16. In the **Messages** dialog box, click **Save**.
17. In the **Save to Master Configuration** dialog box, complete one of the following steps:
  - If you are under a Network Deployment environment, make sure the check box **Synchronize changes with Nodes** is selected, and then click **Save**.
  - If you are not under a Network Deployment environment, click **Save**.
18. In the navigation pane, click **Environment > WebSphere Variables**.
19. Delete the following variables:
  - **AM\_CONFIG\_JVM\_ARGS**
  - **AM\_OLD\_JVM\_ARGS**
  - **ITCAMDCHOME**
  - **ITCAMDCVERSION**
20. In the **Messages** dialog box, click **Save**.
21. In the **Save to Master Configuration** dialog box, complete one of the following steps:
  - If you are under a Network Deployment environment, make sure the check box **Synchronize changes with Nodes** is selected, and then click **Save**.
  - If you are not under a Network Deployment environment, click **Save**.
22. If you configured the server instance for data collection with the data collector configuration tool, rather than manually, complete the following steps:
  - a) Navigate to the `dc_home/runtime` directory.
  - b) Rename the `$profile.$cell.$node.$server.input.properties` file to `$profile.$cell.$node.$server.input.properties.bak`.
23. If you are manually removing the data collector configuration from all application server instances in a profile, perform the following steps:
  - a) Navigate to the `$appserverhome/bin` directory.
  - b) Run the **osgiCfgInit.bat -all** command on Windows systems or the **osgiCfgInit.sh -all** command on UNIX and Linux systems.
24. Restart the application server instance that was monitored by the data collector.



---

# Chapter 14. Configuring the ICAM Agents

Some of the ICAM Agents require configuration. Review the common procedures and agent-specific topics to learn about the default settings and configuration options.

## Common procedures

After installation, some agents are configured and started automatically, while some agents require manual configuration but start automatically. Some agents must be configured and started manually. Multiple instance agents require creating a first instance and starting manually.

### Before you begin

When you install an agent, a sample silent configuration file is placed in the *install\_dir/samples* directory, for example, */opt/ibm/apm/agent/samples/iib\_silent\_config.txt*.

**Remember:** Some agents, for example, the WebSphere Applications agent, have multiple silent configuration files for different tasks such as configuring the data collector.

### About this task

To configure an agent, you can use the command line or a silent response file as described in this procedure.

Configuration methods vary across agents. Some configuration method might not be supported for your agent on a specific operating system. Use the procedure that is provided for your agent.

For more information about agent commands, see [“Using agent commands” on page 226](#).

### Procedure

- To configure the agent with interaction by responding to prompts, run the following one of the commands:

- For single instance agents, run the following command:

- Linux | UNIX

```
agent-name.sh config
```

- Windows

```
agent-name.bat config
```

- For multiple-instance agents, run the following command:

- Linux | UNIX `agent-name.sh config instance_name`

- Windows `agent-name.bat config instance_name`

where:

- *agent-name* is the name of the agent that is specified in [Table 20 on page 226](#).
- *instance\_name* is the instance name, which can be assigned to indicate what you are monitoring.
- To configure the agent without interaction, edit the silent response file and then run one of the following commands:
  - For single instance agents, run the following command:

- Linux | UNIX

```
agent-name.sh config response_file
```

- **Windows**

```
agent-name.bat config response_file
```

- For multiple-instance agents, run the following command:

- **Linux** | **UNIX** `agent-name.sh config instance_name response_file`

- **Windows** `agent-name.bat config instance_name response_file`

where:

- `agent-name` is the name of the agent that is specified in [Table 20](#) on page 226.
- `instance_name` is the instance name, which can be assigned to indicate what you are monitoring.
- `response_file` is the path to the silent response file.

## Using agent commands

The same scripts that you use to install monitoring agents are also used to check the status of an installed agent, stop or start it, or uninstall the agent.

### About this task

The agent name and agent codes are provided for your reference.

Use the agent name in the following commands:

- **Linux** | **UNIX**

```
name-agent.sh
```

- **Windows**

```
name-agent.bat
```

where `name` is the name of the agent that is specified in [Table 20](#) on page 226.

Monitoring agent	<i>name</i>	Two letter agent code
Amazon EC2 agent	amazon_ec2	b5
Amazon ELB agent	amazon_elb	a1
Azure Compute agent	azure_compute	ak
Cassandra agent	cassandra	zc
Cisco UCS agent	ciscoucs	v6
Citrix VDI agent	citrix_vdi	vd
CouchDB agent	couchdb	ck
DataPower agent	datapower	bn
Db2 agent	db2	ud
DataStage agent	datastage	td
Hadoop agent	hadoop	h8
Microsoft Hyper-V Server agent	hyper-v	hv

Table 20. Agent names and agent codes (continued)

<b>Monitoring agent</b>	<b>name</b>	<b>Two letter agent code</b>
IBM Integration Bus agent	iib	qi
JBoss agent	jboss	je
Linux OS agent	os	lz
Linux KVM agent	linux_kvm	v1
MariaDB agent	mariadb	mj
Microsoft .NET agent	dotnet	qe
Microsoft Active Directory agent	msad	3z
Microsoft Cluster Server agent	mscs	q5
Microsoft Exchange Server agent	msexch	ex
Microsoft IIS agent	msiis	q7
Microsoft Office 365 agent	microsoft_office365	mo
Microsoft SQL Server agent	mssql	oq
MongoDB agent	mongodb	ki
MySQL agent	mysql	se
NetApp Storage agent	netapp	nu
Oracle Database agent	oracle_database	rz
PostgreSQL agent	postgresql	pn
RabbitMQ agent	rabbitMQ	zr
SAP agent	sap	sa
SAP HANA Database agent	sap_hana	s7
SAP NetWeaver Java Stack agent	sap_netweaver_java_stack	sv
Skype for Business Server agent	skype for business server	ql
Sterling Connect Direct agent	sterling connect direct	fc
Sterling File Gateway agent	sterling file gateway	fg
Sybase agent	sybase	oy
Tomcat agent	tomcat	ot
UNIX OS agent	os	ux
VMware VI agent	vmware_vi	vm
WebLogic agent	weblogic	wb
WebSphere Applications agent	was	yn
WebSphere Infrastructure Manager agent	wim	d0

Table 20. Agent names and agent codes (continued)

Monitoring agent	name	Two letter agent code
IBM MQ(formerly WebSphere MQ) agent	mq	mq
Windows OS agent	os	nt

## Procedure

- Linux UNIX

On the system where you want to send a command to the monitoring agent, change to the *install\_dir/bin* directory, for example, */opt/ibm/apm/agent/bin*. Enter any of the commands in Table 21 on page 228, where *name* is the agent name that is specified in Table 20 on page 226.

Table 21. Commands for AIX and Linux systems

Command	Description
<code>./name-agent.sh status</code>	Checks the agent status. Status can be either running or not running. When the agent is running, the connection status between the agent and the Cloud App Management server is also checked. Possible negative connection statuses are: Connection failed, Error detected, and Disconnected-error. The positive status is Connected, this is the expected status. The transitional status is Connecting. A status of Unknown means that the agent status cannot be recognized, which is possibly due to errors in the file system or in the agent log file.
<code>./name-agent.sh start</code>	Starts the monitoring agent. If the agent has instances, enter an instance name after the command.
<code>./name-agent.sh stop</code>	Stops the agent. If the agent has instances, enter an instance name after the command.
<code>./name-agent.sh prereqcheck</code>	Runs a prerequisite scan. This command option is available for most agents.
<code>./name-agent.sh install</code>	Installs the monitoring agent.
<code>./name-agent.sh config instance_name path_to_silent_config_file</code>	Configures the monitoring agent. Run the command from the <i>install_dir/bin</i> directory and add the response file path if required.  If the agent has instances, enter an instance name. For more information about which agents are multiple instance agents, see the Table 18 on page 210.  The <i>silent_config_file</i> is optional. If you do not specify a file for silent configuration, you can configure the monitoring agent interactively by following the prompts.
<code>./name-agent.sh uninstall</code>	Uninstalls the monitoring agent. For more information, see “Uninstalling your agents” on page 214.



Table 21. Commands for AIX and Linux systems (continued)	
Command	Description
<code>./smal-agent.sh uninstall_all</code>	Uninstalls all the monitoring agents on the managed system.
<code>./name-agent.sh remove instance_name</code>	Removes an instance of a multiple instance agent.
<code>./name-agent.sh</code>	View a description of the functions that are available with the script.

#### Windows

On the system or VM where you want to send a command to the monitoring agent, change to the `install_dir\BIN` directory at the command prompt, for example, `C:\IBM\APM\BIN`. Enter any of the commands in [Table 22 on page 229](#), where `name` is the agent name that is specified in [Table 20 on page 226](#).

Table 22. Commands for Windows systems	
Command	Description
<code>name-agent.bat status</code>	Checks the agent status.  Checks the connection status between the agent and the Cloud App Management server. Possible negative connection statuses are: Connection failed, Error detected, and Disconnected-error. The positive status is Connected, this is the expected status. The transitional status is Connecting. A status of Unknown means that the agent status cannot be recognized, which is possibly due to errors in the file system or in the agent log file.
<code>name-agent.bat start</code>	Starts the monitoring agent. If the agent has instances, enter an instance name after the command.
<code>name-agent.bat stop</code>	Stops the agent. If the agent has instances, enter an instance name after the command.
<code>name-agent.bat prereqcheck</code>	Runs a prerequisite scan. This command option is available for most agents.
<code>name-agent.bat install</code>	Installs the monitoring agent.
<code>name-agent.bat config instance_name path_to_silent_config_file</code>	Configures the monitoring agent. Run the command from <code>install_dir\BIN</code> directory and add the response file path if required.  If the agent has instances, enter an instance name. For more information about which agents are multiple instance agents, see the <a href="#">Table 18 on page 210</a> .  The <code>silent_config_file</code> is optional. If you do not specify a file for silent configuration, you can configure the monitoring agent interactively by following the prompts.

Table 22. Commands for Windows systems (continued)	
Command	Description
<code>name-agent.bat uninstall</code>	Uninstalls the monitoring agent. For more information, see <a href="#">“Uninstalling your agents” on page 214.</a>
<code>smai-agent.bat uninstall_all</code>	Uninstalls all monitoring agents on the managed system.
<code>name-agent.bat remove instance_name</code>	Removes an instance of a multiple instance agent.
<code>name-agent.bat</code>	View a description of the functions that are available with the script.

Agent version command

- To see the version of an agent in your environment, run the following commands:

– **Linux** | **UNIX**

```
install_dir/bin/cinfo
```

– **Windows**

```
install_dir\InstallITM\kincinfo
```

## Starting agents as a non-root user

If you want to start agents as different users, create a common group on the system and make each user a member of this group.

### Before you begin

If you installed and configured your agent as the same non-root user and you want to start the agent as the same user, no special action is required.

If you installed and configured your agent as a selected user and want to start the agent as a different user, create a common group on the system. Make all agent management users members of this common group. Transfer ownership of all agent files and directories to this group.

### About this task

An autostart script is generated by an agent installation, upgrade, or configuration. This script (named ITMAgentsN or rc.itmN, depending on the UNIX operating system) contains an entry for each application in a particular installation. By default all agents are started with root user access.

To update system startup scripts and start agents as a non-root user, you must edit the `install_dir/config/kcirunas.cfg` file, which contains a superset of the XML syntax.

Each **productCode** section in the `kcirunas.cfg` file is disabled by default. Activate a **productCode** section for your agent by removing the comment indicator from **!productCode**. Commented or deactivated sections are ignored. Uncommented or activated sections for applications that are not installed are ignored.

### Procedure

1. Install your monitoring agents on Linux or AIX as described in [“Installing agents” on page 198](#) on AIX systems or [“Installing agents” on page 202](#) on Linux systems.
2. Optional: Configure your monitoring agents on Linux or AIX as necessary, see [Chapter 14, “Configuring the ICAM Agents,” on page 225.](#)

3. Run the following command from the `install_dir/bin` directory with the group name of the non-root user to secure the files and set the file group ownership to the files.

```
./secure.sh -g group_name
```

For example:

```
./secure.sh -g mqadmin1
```

4. To update the system startup scripts, complete the following steps:

- a) Update the `install_dir/config/kcirusnas.cfg` file. Activate **productCode** sections for your agents.

For agents that do not require an instance value, specify the **product\_code** and **user** values, where the `product_code` value is the two-letter code that is specified in [Table 20 on page 226](#). For agents that do require an instance value, such as the IBM MQ(formerly WebSphere MQ) agent (product code: mq), specify the **product\_code**, **user**, and **name** values, where **name** is the instance name.

For example:

```
<productCode>mq</productCode>
<instance>
<name>qmgrinst1</name>
<user>qmgrinst1</user>
</instance>
<instance>
<name>qmgrinst2</name>
<user>root</user>
</instance>
```

- b) Run the following command with root user or sudo user access:

```
install_dir/bin/UpdateAutoRun.sh
```

## Results

The agents can be started by a non-root user, which is not the same user that installed and configured the agents. You can use the same user ID for agent upgrades.

For more information about the `./secure.sh` script, see [Securing the agent installation files](#).

## Configuring agents as a non-root user

If you want to configure your agent as a non-root user, create a common group on the system and make each user a member of this group.

### Before you begin

If you installed your agent as a root or non-root user and you want to configure the agent as the same user, no special action is required.

If you installed your agent as a selected user and want to configure the agent as a different user, create a common group on the system. Make all agent management users members of this common group. Transfer ownership of all agent files and directories to this group.

**Remember:** For the IBM Integration Bus agent, if IBM Integration Bus installation is a single-user deployment, use the same user ID as the user who installed IBM Integration Bus to configure the agent. Before you configure the agent, complete the following steps for this user ID.

### Procedure

1. Install your monitoring agents on Linux or AIX as described in [“Installing agents on Linux systems” on page 200](#) and [“Installing agents on UNIX systems” on page 196](#).

2. Run the following command from the `install_dir/bin` directory with the group name of the non-root user to secure the files and set the file group ownership to the files.

```
./secure.sh -g group_name
```

For example:

```
./secure.sh -g mqadmin1
```

3. Configure your monitoring agents on Linux or AIX as necessary, see [Chapter 14, “Configuring the ICAM Agents,”](#) on page 225.
4. To update the system startup scripts, run the following script with root user or sudo user access:

```
install_dir/bin/UpdateAutoRun.sh
```

## Results

The agents can be configured by the non-root user, which is not the same user that installed and configured the agents. You can use the same user ID for agent installation and upgrades.

For more information about the `./secure.sh` script, see [Securing the agent installation files](#).

## Disabling automatic agent start on AIX and Linux systems

On an AIX or Linux system, an agent can automatically start after operating system restart. If you do not want the agent to start automatically after system restart, you can disable automatic agent start.

### About this task

If you install an agent as root user on the AIX or Linux system, the agent can automatically start after system restart. Or, if you install an agent as non-root user but run the `UpdateAutoRun.sh` script as root after installation, the agent can automatically start after system restart.

### Procedure

To disable automatic agent start, complete the following steps:

1. Open the `install_dir/registry/AutoStart` file in a text editor of your choice.
2. Change the content to 0 and save your changes.

The previous content is a positive number, such as 1, 2, 3, or 4, which specifies the agent script to run after system restart.

### Results

After system restart, no agent script will automatically run to start the agent.

## Configuring Amazon EC2 monitoring

The Monitoring Agent for Amazon EC2 offers a central point of management for your Amazon EC2 environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Amazon EC2 agent you can easily collect and analyze Amazon EC2 specific information.

### Before you begin

- Read the entire [“Configuring Amazon EC2 monitoring”](#) on page 232 topic to determine what is needed to complete the configuration.
- Make sure that the system requirements for the Amazon EC2 agent are met in your environment. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\)](#) for the Amazon EC2 agent.

- Ensure that the following information is available:
  - A list of the AWS region names that contain EC2 instances to monitor.
  - The AWS security credentials (Access Key ID and Secret Access Key) with permission to access each AWS region.
- Ensure that the AWS security credentials that are used for each AWS region are a member of a group that includes at least the *AmazonEC2ReadOnlyAccess* policy.

### About this task

The Amazon EC2 agent is both a multiple instance agent and also a subnode agent. You can create one agent instance with multiple subnodes – one for each Amazon EC2 region, or you can create an agent instance for each Amazon EC2 region with one subnode for that region. Or you can create a combination of each type of configuration. After you configure agent instances, you must start each agent instance manually. If you have more than 50 resources per Amazon EC2 region, it is suggested that you create an agent instance per region or use tagging on your EC2 instances and filter agent instances by the tags you create by using the agent's [filtering condition](#) parameter.

### Procedure

1. Configure the agent on Windows systems with the **IBM Performance Management** window or the silent response file.
  - [“Configuring the agent on Windows systems”](#) on page 235.
  - [“Configuring the agent by using the silent response file”](#) on page 237.
2. Configure the agent on Linux systems with the script that prompts for responses or the silent response file.
  - [“Configuring the agent by responding to prompts”](#) on page 236.
  - [“Configuring the agent by using the silent response file”](#) on page 237.

### What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI”](#) on page 176.

If you are unable to view the data in the agent dashboards, first check the server connection logs and then the data provider logs. The default paths to these logs are listed here:

- **Linux** /opt/ibm/apm/agent/logs
- **Windows** C:\IBM\APM\TMAITM6\_x64\logs

## Configuration parameters for the Amazon EC2 agent

The configuration parameters for the Amazon EC2 agent are displayed in a table.

1. [Amazon EC2 Region Configuration](#) - Settings to monitor Amazon EC2 instances remotely. Instances are automatically discovered in the specified region that you want to configure.

Table 23. Amazon EC2 Region Configuration

Parameter name	Description	Silent configuration file parameter name
EC2 Subnode Name	<p>Name of the EC2 Subnode for collection of data. Example, <i>usw2a</i>.</p> <p>This alias is part of the managed system name (MSN) and it is used to visually identify the monitored environment in the Cloud APM console.</p> <p><b>Note:</b> This alias can be anything that you choose to represent the Amazon EC2 subnode instance with the following restrictions. Letters from the Latin alphabet (a-z, A-Z), Arabic numerals (0-9), the hyphen-minus character (-), and the underscore character (_) can be used to create agent subnode instance names. The maximum length of an EC2 subnode name is 25 characters.</p>	<p>Each of the following parameters must have an agent subnode name suffix that is the same for each parameter of an agent subnode instance. New agent subnode instances must use a unique name for its set of parameters. For example, one agent subnode instance can use <i>.usw2a</i> and another agent subnode instance can use <i>.usw2b</i> in place of <i>.subnode_name</i> in the parameter names that follow.</p>
Access ID	<p>AWS security credentials Access Key ID that is used to authenticate with the specified Amazon Region. For example, 'AKIAxxxxxxxxxxxxxxxx'.</p>	<p><b>KB5_INS_ACCESS_ID.subnode_name</b></p>
Secret Key	<p>AWS security credentials Secret Access Key that is used to authenticate with the specified Amazon Region. For example, 'kK7txxxxxxxxxxxxxxxxxxxxxxxxxxxx'.</p>	<p><b>KB5_INS_SECRET_KEY.subnode_name</b></p>
Region	<p>AWS Region to monitor. For example, 'us-west-2'.</p>	<p><b>KB5_INS_REGION.subnode_name</b></p>

Table 23. Amazon EC2 Region Configuration (continued)

Parameter name	Description	Silent configuration file parameter name
Filtering Condition	<p>The type of filtering that is being done.</p> <p>You can use custom tags on EC2 instances to limit which EC2 instances are monitored by the agent. For more information, see <a href="#">Tagging Your Amazon EC2 Resources</a>.</p> <p>Filtering options,</p> <p><b>none</b> All EC2 instances within the region are monitored. <b>Filter Value</b> is ignored.</p> <p><b>tagName</b> EC2 instances with the tag key that is specified in <b>Filter Value</b> are monitored, regardless of the actual value in the corresponding EC2 instance tag value. For example, to monitor all EC2 instances that have the tag key <i>Stack</i>, regardless of the value in its tag value, specify <i>Stack</i> in <b>Filter Value</b>.</p> <p><b>tagName tagValue</b> EC2 instances with the tag key and tag value pair that is separated with a vertical bar ( ), and that is specified in <b>Filter Value</b> are monitored. For example, to monitor all EC2 instances that have the tag key <i>Stack</i> and the tag value <i>Production</i>, specify <i>Stack Production</i> in <b>Filter Value</b>.</p> <p><b>monitoring-tag</b> EC2 instances that have at least one tag are monitored. <b>Filter Value</b> is ignored.</p>	<p><b>FILTER_CONDITION.subnode_name</b></p> <p>Valid values,</p> <p><b>none</b> none</p> <p><b>tagName</b> tagName</p> <p><b>tagName tagValue</b> tagName tagValue</p> <p><b>monitoring-tag</b> monitoring-tag</p>
Filter Value	<p>The value of the tag by which the EC2 instances are filtered when either <b>tagName</b> or <b>tagName tagValue</b> are selected for <b>Filtering Condition</b>.</p>	<p><b>FILTER_VALUE.subnode_name</b></p>

## Configuring the agent on Windows systems

You can configure the Amazon EC2 agent on Windows operating systems by using the IBM Cloud App Management window. After you update the configuration values, you must start the agent to save the updated values.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Cloud App Management**.
2. In the **IBM Performance Management** window, right-click the **Monitoring Agent for Amazon EC2** template, and then click **Configure agent**.

**Remember:** After you configure an agent instance for the first time, the **Configure agent** option is disabled. To configure the agent instance again, right-click on it and then click **Reconfigure...**

3. Enter a unique instance name then click **OK**. Use only Latin letters, Arabic numerals, and the hyphen-minus character in the instance name. Example, `ec2-inst3`.
4. Click **Next** on the agent instance name window.
5. Enter the **Amazon EC2 Region Configuration** instance template settings.

**Note:** This section is not the Amazon EC2 region instance configuration. It is a template section for setting what is used as the default values when you add the actual Amazon EC2 region instance configurations in [step 6](#).

See [Table 23 on page 234](#) for an explanation of each of the configuration parameters.

6. Press **New** and enter Amazon EC2 region instance settings, then click **Next**.  
See [Table 23 on page 234](#) for an explanation of each of the configuration parameters.
7. Click **OK** to complete the configuration.
8. In the IBM Cloud App Management window, right-click the instance that you configured, and then click **Start**.

## Configuring the agent by responding to prompts

After installation of the Amazon EC2 agent, you must configure it before you start the agent. If the Amazon EC2 agent is installed on a local Linux computer, you can follow these instructions to configure it interactively through command line prompts.

### About this task

**Remember:** If you are reconfiguring a configured agent instance, the value that is set in the last configuration is displayed for each setting. If you want to clear an existing value, press the space key when the setting is displayed.

### Procedure

Follow these steps to configure the Amazon EC2 agent by running a script and responding to prompts.

1. Run the following command:

```
install_dir/bin/amazonec2-agent.sh config instance_name
```

where `install_dir` is the path where the agent is installed and `instance_name` is the name that you want to give to the agent instance.

Example

```
/opt/ibm/apm/agent/bin/amazonec2-agent.sh config ec2-inst3
```

2. Respond to the prompts to set configuration values for the agent.

See [“Configuration parameters for the Amazon EC2 agent” on page 233](#) for an explanation of each of the configuration parameters.

3. Run the following command to start the agent:

```
install_dir/bin/amazonec2-agent.sh start instance_name
```

where `install_dir` is the path where the agent is installed and `instance_name` is the name of the agent instance.

Example

```
/opt/ibm/apm/agent/bin/amazonec2-agent.sh start ec2-inst3
```



## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

The silent response file contains the agent configuration parameters with default values that are defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters.

After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

- Configure the Amazon EC2 agent in the silent mode:

a) Open the `amazonec2_silent_config.txt` file at one of the following paths in a text editor.

– **Linux** `install_dir/samples/amazonec2_silent_config.txt`  
Example, `/opt/ibm/apm/agent/samples/amazonec2_silent_config.txt`

– **Windows** `install_dir\samples\amazonec2_silent_config.txt`  
Example, `C:\IBM\APM\samples\amazonec2_silent_config.txt`

where `install_dir` is the path where the agent is installed.

b) In the `amazonec2_silent_config.txt` file, specify values for all mandatory parameters and modify the default values of other parameters as needed.

See “Configuration parameters for the Amazon EC2 agent” on page 233 for an explanation of each of the configuration parameters.

c) Save and close the `amazonec2_silent_config.txt` file, and run the following command:

– **Linux** `install_dir/bin/amazonec2-agent.sh config instance_name install_dir/samples/amazonec2_silent_config.txt`

Example, `/opt/ibm/apm/agent/bin/amazonec2-agent.sh config ec2-inst3 /opt/ibm/apm/agent/samples/amazonec2_silent_config.txt`

– **Windows** `install_dir\bin\amazonec2-agent.bat config instance_name install_dir\samples\amazonec2_silent_config.txt`

Example, `C:\IBM\APM\bin\amazonec2-agent.bat config ec2-inst3 C:\IBM\APM\samples\amazonec2_silent_config.txt`

where `install_dir` is the path where the agent is installed and `instance_name` is the name that you want to give to the agent instance.

**Important:** Ensure that you include the absolute path to the silent response file. Otherwise, the agent data is not shown in the dashboards.

d) Run the following command to start the agent:

– **Linux** `install_dir/bin/amazonec2-agent.sh start instance_name`

Example, `/opt/ibm/apm/agent/bin/amazonec2-agent.sh start ec2-inst3`

– **Windows** `install_dir\bin\amazonec2-agent.bat start instance_name`

Example, `C:\IBM\APM\bin\amazonec2-agent.bat start ec2-inst3`

where `install_dir` is the path where the agent is installed and `instance_name` is the name of the agent instance.

### What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring AWS Elastic Load Balancer monitoring

The Monitoring Agent for AWS Elastic Load Balancer offers a central point of management for your AWS Elastic Load Balancer environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Amazon ELB agent you can easily collect and analyze AWS Elastic Load Balancer specific information.

### Before you begin

- Read the entire [“Configuring AWS Elastic Load Balancer monitoring” on page 238](#) topic to determine what is needed to complete the configuration.
- Make sure that the system requirements for the Amazon ELB agent are met in your environment. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\) for the Amazon ELB agent](#).
- Ensure that the following information is available:
  - The AWS security credentials (Access Key ID and Secret Access Key) with permission to access each AWS region with Elastic Load Balancers.

### About this task

The Amazon ELB agent is both a multiple instance agent and also a subnode agent. The subnodes are created automatically for each type of Elastic Load Balancer that is available in your AWS environment.

### Procedure

1. Configure the agent on Windows systems with the **IBM Performance Management** window or the silent response file.
  - [“Configuring the agent on Windows systems” on page 239](#).
  - [“Configuring the agent by using the silent response file” on page 241](#).
2. Configure the agent on Linux systems with the script that prompts for responses or the silent response file.
  - [“Configuring the agent by responding to prompts” on page 239](#).
  - [“Configuring the agent by using the silent response file” on page 241](#).

### What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

If you are unable to view the data in the agent dashboards, first check the server connection logs and then the data provider logs. The default paths to these logs are listed here:

- **Linux** /opt/ibm/apm/agent/logs
- **Windows** C:\IBM\APM\TMAITM6\_x64\logs

## Configuration parameters for the Amazon ELB agent

The configuration parameters for the Amazon ELB agent are displayed in a table.

1. [Table 24 on page 239](#) - Credentials that are required for access to the Amazon Subscription that contains the AWS Elastic Load Balancers to monitor.

Table 24. Subscription information

Parameter name	Description	Silent configuration file parameter name
Access Key ID	The access ID that is used to authenticate with the specified Amazon Region. For example, 'AKIAxxxxxxxxxxxxxxxx'.	<b>KAL_ACCESS_KEY_ID</b>
Secret Access Key	The secret access key that is used to authenticate with the specified Amazon Region. For example, 'kK7txxxxxxxxxxxxxxxxxxxxxxxxxxxx'.	<b>KAL_SECRET_ACCESS_KEY_PASSWORD</b>
Region	The Amazon region where the load balancers are located. For example, 'us-west-2'.	<b>KAL_REGION</b>

## Configuring the agent on Windows systems

You can configure the Amazon ELB agent on Windows operating systems by using the IBM Cloud App Management window. After you update the configuration values, you must start the agent to save the updated values.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Cloud App Management**.
2. In the **IBM Performance Management** window, right-click the **Monitoring Agent for AWS Elastic Load Balancer** template, and then click **Configure agent**.
 

**Remember:** After you configure an agent instance for the first time, the **Configure agent** option is disabled. To configure the agent instance again, right-click on it and then click **Reconfigure...**
3. In the Monitoring Agent for AWS Elastic Load Balancer window, complete the following steps:
  - a) Enter a unique instance name for the Monitoring Agent for AWS Elastic Load Balancer instance, and click **OK**.
4. Enter the **Amazon ELB Subscription Credentials**, then click **Next**.
 

See “Configuration parameters for the Amazon ELB agent” on page 238 for an explanation of each of the configuration parameters. If your **Secret Key/Password** contains an equal sign (=), you must reenter it each time that you reconfigure the agent.
5. Click **OK** to complete the configuration.
6. In the IBM Cloud App Management window, right-click the instance that you configured, and then click **Start**.

### What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent by responding to prompts

After installation of the Amazon ELB agent, you must configure it before you start the agent. If the Amazon ELB agent is installed on a local Linux computer, you can follow these instructions to configure it interactively through command line prompts.

### About this task

**Remember:** If you are reconfiguring a configured agent instance, the value that is set in the last configuration is displayed for each setting. If you want to clear an existing value, press the space key when the setting is displayed.

## Procedure

Follow these steps to configure the Amazon ELB agent by running a script and responding to prompts.

1. Run the following command:

```
install_dir/bin/amazon_elb-agent.sh config instance-name
```

Where *install\_dir* is the path where the agent is installed and *instance-name* is the name that you want to give to the agent instance.

Example

```
/opt/ibm/apm/agent/bin/amazon_elb-agent.sh config elb-inst3
```

2. Respond to the prompts to set configuration values for the agent.

See [“Configuration parameters for the Amazon ELB agent” on page 238](#) for an explanation of each of the configuration parameters.

3. Run the following command to start the agent:

```
install_dir/bin/amazon_elb-agent.sh start instance-name
```

Where *install\_dir* is the path where the agent is installed and *instance-name* is the name of the agent instance.

Example

```
/opt/ibm/apm/agent/bin/amazon_elb-agent.sh start elb-inst3
```

## Example

Creating an agent instance that is named `elb-inst3`.

```
# ./amazon_elb-agent.sh config elb-inst3
Configuring Monitoring Agent for Amazon ELB

Edit 'Monitoring Agent for Amazon ELB' settings? [ 1=Yes, 2=No ] (default is: 1): 1

Subscription Information :
Amazon ELB subscription information

The access ID that is used to authenticate with the specified Amazon Region.
For example, 'AKIAxxxxxxxxxxxxxxxx'.
Access Key ID (default is: ): AKIAIOSFODNN7EXAMPLE

The secret access key that is used to authenticate with the specified Amazon
Region. For example, 'kK7txxxxxxxxxxxxxxxxxxxxxxxxxxxx'.
Enter Secret Access Key (default is: ): hidden

Re-type : Secret Access Key (default is: ): hidden

The Amazon region where the load balancers are located. For example, 'us-west-2'.
Region (default is: ): us-west-2

Configuration completed successfully.
Automatic start at system initialization has been configured.
Automatic stop at system shutdown has been configured.
```

## What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

The silent response file contains the agent configuration parameters with default values that are defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters.

After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

Follow these steps to configure the Amazon ELB agent in the silent mode.

1. Open the `amazon_elb_silent_config.txt` file at one of the following paths in a text editor.

- **Linux** `install_dir/samples/amazon_elb_silent_config.txt`  
Example, `/opt/ibm/apm/agent/samples/amazon_elb_silent_config.txt`
- **Windows** `install_dir\samples\amazon_elb_silent_config.txt`  
Example, `C:\IBM\APM\samples\amazon_elb_silent_config.txt`

Where `install_dir` is the path where the agent is installed.

2. In the `amazon_elb_silent_config.txt` file, specify values for all mandatory parameters and modify the default values of other parameters as needed.

See “[Configuration parameters for the Amazon ELB agent](#)” on page 238 for an explanation of each of the configuration parameters.

3. Save and close the `amazon_elb_silent_config.txt` file, and run the following command:

- **Linux** `install_dir/bin/amazon_elb-agent.sh config instance-name install_dir/samples/amazon_elb_silent_config.txt`  
Example, `/opt/ibm/apm/agent/bin/amazon_elb-agent.sh config elb-inst3 /opt/ibm/apm/agent/samples/amazon_elb_silent_config.txt`
- **Windows** `install_dir\bin\amazon_elb-agent.bat config instance-name install_dir\samples\amazon_elb_silent_config.txt`  
Example, `C:\IBM\APM\bin\amazon_elb-agent.bat config elb-inst3 C:\IBM\APM\samples\amazon_elb_silent_config.txt`

Where `install_dir` is the path where the agent is installed and `instance-name` is the name that you want to give to the agent instance.

4. Run the following command to start the agent:

- **Linux** `install_dir/bin/amazon_elb-agent.sh start instance-name`  
Example, `/opt/ibm/apm/agent/bin/amazon_elb-agent.sh start elb-inst3`
- **Windows** `install_dir\bin\amazon_elb-agent.bat start instance-name`  
Example, `C:\IBM\APM\bin\amazon_elb-agent.bat start elb-inst3`

Where `install_dir` is the path where the agent is installed and `instance-name` is the name of the agent instance.

## Example

Edited amazon\_elb\_silent\_config.txt.

```
#
# This is a sample configuration response file for agent Amazon ELB.
#
# It contains an entry for every configuration property.
# Entries for optional properties that have no default value are included
# in comments.
# Ensure that all uncommented properties have a value before configuring
# the agent.
#

# Access Key ID: The access ID that is used to authenticate with the
# specified Amazon Region. For example, 'AKIAxxxxxxxxxxxxxxxx'.
KAL_ACCESS_KEY_ID=AKIAIOSFODNN7EXAMPLE

# Secret Access Key: The secret access key that is used to authenticate with
# the specified Amazon Region. For example, 'kK7txxxxxxxxxxxxxxxxxxxxxxxx'.
KAL_SECRET_ACCESS_KEY_PASSWORD=wJa1rXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY

# Region: The Amazon region where the load balancers are located. For
# example, 'us-west-2'.
KAL_REGION=us-west-2
```

## What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring Azure Compute monitoring

The Monitoring Agent for Azure Compute offers a central point of management for your Azure Compute environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the Azure Compute agent you can easily collect and analyze Azure Compute specific information.

### Before you begin

- Read the entire [“Configuring Azure Compute monitoring” on page 242](#) topic to determine what is needed to complete the configuration.
- Make sure that the system requirements for the Azure Compute agent are met in your environment. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\) for the Azure Compute agent](#).
- Ensure that the following information is available:
  - The Azure Subscription Credentials with permission to access the Azure Compute instances to monitor. See [“Azure Compute Configuration Information” on page 243](#) for more details.

### About this task

The Azure Compute agent is both a multiple instance agent and also a subnode agent. Each Azure Compute agent subnode monitors a grouping of Azure Compute virtual machines according to a filter you define. You can create one agent instance with multiple subnodes – one for each virtual machine grouping, or you can create an agent instance for each virtual machine grouping with one subnode for that grouping. Or you can create a combination of each type of configuration. After you configure agent instances, you must start each agent instance manually. It is suggested that you have no more than 50 resources per Azure Compute virtual machine grouping. Each Azure Compute agent subnode name must be unique within your environment.

## Procedure

1. Configure the agent on Windows systems with the **IBM Performance Management** window or the silent response file.
  - [“Configuring the agent on Windows systems” on page 246.](#)
  - [“Configuring the agent by using the silent response file” on page 248.](#)
2. Configure the agent on Linux systems with the script that prompts for responses or the silent response file.
  - [“Configuring the agent by responding to prompts” on page 247.](#)
  - [“Configuring the agent by using the silent response file” on page 248.](#)

## What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176.](#)

If you are unable to view the data in the agent dashboards, first check the server connection logs and then the data provider logs. The default paths to these logs are listed here:

- **Linux** /opt/ibm/apm/agent/logs
- **Windows** C:\IBM\APM\TMAITM6\_x64\logs

## Azure Compute Configuration Information

The Azure Compute agent requires the some additional setup in the Azure Compute environment.

### About this task

To run these steps, you must login to the [Microsoft Azure console](#).

## Procedure

1. Subscription ID
  - On the left pane, select "Subscriptions" and chose the subscription you want to use for this agent.
  - Select "Overview", then copy the Subscription ID. This will be used as one of the Agent's configuration parameters.
2. Tenant ID
  - Navigate to "Azure Active Directory".
  - Select "Properties", then copy the Tenant ID.
3. Register an Application
  - Go to "All services" and search for "App registrations".
  - Click "New Application Registration".
  - Fill out a name, select Application Type "Web App/API", and a sign-on URL (this URL will not be used so chose whatever you'd like).
  - Click "Create"
  - Copy the Application ID - This will be used in the Agent's "Client ID" field.
4. Create Application Key
  - Click on the App you just created, then go to "Settings" followed by "Keys".
  - Enter a description (e.g., "IBM Key") and duration (e.g., "Never Expires"), then click "Save".
  - Copy the Secret Key and store it somewhere safe - you will only see this key one time and will need to generate a new one if you lose it.
5. Give the Application Permissions

- Go to "Subscriptions" and select the subscription to be monitored.
- Go to "Access Control (IAM) and click "Add".
- Select "Reader" role or higher for monitoring.
- Under "Select", find the App you just registered and select it, then click "Save".

## Configuration parameters for the Azure Compute agent

The configuration parameters for the Azure Compute agent are displayed in tables which group them according to sections.

1. [Table 25 on page 244](#) - Credentials that are required for access to the Azure Subscription that contains the Azure Compute resources to monitor.
2. [Table 26 on page 245](#) - Create agent subnodes to define groupings of virtual machines. Each subnode name must be unique within your environment. It is suggested that you have no more than 50 virtual machines per subnode.

Parameter name	Description	Silent configuration file parameter name
Subscription ID	The ID assigned by Azure for the Subscription that is monitored.	<b>KAK_SUBSCRIPTION_ID</b>
Tenant ID	The tenant ID that is assigned by Azure. Used to log in to the Azure service API.	<b>KAK_TENANT_ID</b>
Client ID	The client ID that is assigned by Azure to identify this agent as an external application that monitors the Azure compute services.	<b>KAK_CLIENT_ID</b>
Secret Key/ Password	The secret access key or password that is created by Azure for the client application.  <b>Important:</b> <b>Windows</b> If your <b>Secret Key/Password</b> contains an equal sign (=), you must reenter it each time that you reconfigure the agent.	<b>KAK_SECRET_PASSWORD</b>



Table 26. Azure Compute Virtual Machine Subnode

Parameter name	Description	Silent configuration file parameter name
Subnode Name	<p>Name of the Azure Compute Subnode for collection of data. Example, <i>azc1</i>. Subnode name must be unique in your environment.</p> <p>This alias is part of the managed system name (MSN) and it is used to visually identify the monitored environment in the Cloud APM console.</p> <p><b>Note:</b> This alias can be anything that you choose to represent the Azure Compute subnode instance with the following restrictions. Letters from the Latin alphabet (a-z, A-Z), Arabic numerals (0-9), the hyphen-minus character (-), and the underscore character (_) can be used to create agent subnode instance names. The maximum length of an Azure Compute subnode name is 25 characters.</p>	<p>Each of the following parameters must use a period (.) followed by an agent <b>Subnode Name</b> as a suffix. The <b>Subnode Name</b> must be the same for each subnode parameter. New agent subnode instances must use a unique <b>Subnode Name</b> for its set of parameters. For example, one agent subnode instance can use <i>.azc1</i> and another agent subnode instance can use <i>.azc2</i> in place of <i>.subnode_name</i> in the parameter names that follow.</p>
Filter Type	The type of filter to be applied.	<p><b>KAK_FILTER_TYPE.subnode_name</b></p> <p>Valid values,</p> <p><b>ALL</b> All</p> <p><b>TAG_NAME_VALUE</b> Tag Name-Value Pair</p> <p><b>RESOURCE_GROUP</b> Resource Group</p>

Table 26. Azure Compute Virtual Machine Subnode (continued)

Parameter name	Description	Silent configuration file parameter name
Filter Value	<p>The filter value corresponding to the selected <b>Filter Type</b>. This value can be a <b>Resource Group</b> or <b>Tag Name-Value Pair</b>. Leave it empty for <b>Filter Type All</b>. For command line configuration, a backslash might appear in the example displayed. Do not enter a backslash in the value you provide.</p> <p>Examples of filter type and filter value pairs:</p> <ul style="list-style-type: none"> <li>• Azure Compute Subnode to monitor all virtual machines. Leave the filter value empty. Filter value is not needed and it is ignored for filter type <b>All</b>. <ul style="list-style-type: none"> <li>– Filter Type: <b>All</b></li> <li>– Filter Value:</li> </ul> </li> <li>• Azure Compute Subnode to monitor all virtual machines with a tag name of DTAP and a tag value that matches the string prod. <ul style="list-style-type: none"> <li>– Filter Type: <b>Tag Name-Value Pair</b></li> <li>– Filter Value: DTAP:prod</li> </ul> </li> <li>• Azure Compute Subnode to monitor all virtual machines with a resource group property that matches the string linux-group1. <ul style="list-style-type: none"> <li>– Filter Type: <b>Resource Group</b></li> <li>– Filter Value: linux-group1</li> </ul> </li> </ul>	<b>KAK_FILTER_VALUE.subnode_name</b>

## Configuring the agent on Windows systems

You can configure the Azure Compute agent on Windows operating systems by using the IBM Cloud App Management window. After you update the configuration values, you must start the agent to save the updated values.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Cloud App Management**.
2. In the **IBM Performance Management** window, right-click the **Monitoring Agent for Azure Compute** template, and then click **Configure agent**.

**Remember:** After you configure an agent instance for the first time, the **Configure agent** option is disabled. To configure the agent instance again, right-click on it and then click **Reconfigure....**

3. In the Monitoring Agent for Azure Compute window, complete the following steps:
  - a) Enter a unique instance name for the Monitoring Agent for Azure Compute instance, and click **OK**.
4. Enter the **Azure Subscription Credentials**, then click **Next**.

See [Table 25 on page 244](#) for an explanation of each of the configuration parameters.

**Important:** **Windows** If your **Secret Key/Password** contains an equal sign (=), you must reenter it each time that you reconfigure the agent.

5. Enter the **Azure Compute Virtual Machine Subnode** template settings.

See [Table 26 on page 245](#) for an explanation of each of the configuration parameters.

**Note:** This section is not the Azure Compute Virtual Machine Subnode instance configuration. It is a template section for setting what is used as the default values when you add the actual Azure Compute Virtual Machine Subnode instance configurations in [step 6](#).

6. Press **New** and enter **Azure Compute Virtual Machine Subnode** instance settings, then click **Next**.  
See [Table 26 on page 245](#) for an explanation of each of the configuration parameters.
7. Click **OK** to complete the configuration.
8. In the IBM Cloud App Management window, right-click the instance that you configured, and then click **Start**.

## Configuring the agent by responding to prompts

After installation of the Azure Compute agent, you must configure it before you start the agent. If the Azure Compute agent is installed on a local Linux computer, you can follow these instructions to configure it interactively through command line prompts.

### About this task

**Remember:** If you are reconfiguring a configured agent instance, the value that is set in the last configuration is displayed for each setting. If you want to clear an existing value, press the space key when the setting is displayed.

### Procedure

Follow these steps to configure the Azure Compute agent by running a script and responding to prompts.

1. Run the following command:

```
install_dir/bin/azure_compute-agent.sh config instance-name
```

Where *install\_dir* is the path where the agent is installed and *instance-name* is the name that you want to give to the agent instance. Use only Latin letters, Arabic numerals, and the hyphen-minus character in the *instance-name*.

Example

```
/opt/ibm/apm/agent/bin/azure_compute-agent.sh config azc-inst3
```

2. Respond to the prompts to set configuration values for the agent.

See [“Configuration parameters for the Azure Compute agent” on page 244](#) for an explanation of each of the configuration parameters.

**Remember:** When you first configure an agent instance, you must add at least one subnode when prompted to **Edit 'Azure Compute Virtual Machine Subnode' settings**.

3. Run the following command to start the agent:

```
install_dir/bin/azure_compute-agent.sh start instance-name
```

Where *install\_dir* is the path where the agent is installed and *instance-name* is the name of the agent instance.

Example

```
/opt/ibm/apm/agent/bin/azure_compute-agent.sh start azc-inst3
```

## Example

Creating an agent instance that is named `azc - inst3` and has one subnode instance that is named `azc1`.

```
# ./azure_compute-agent.sh config azc-inst3
Configuring Monitoring Agent for Azure Compute

Edit 'Monitoring Agent for Azure Compute' settings? [ 1=Yes, 2=No ] (default is: 1): 1

Azure Subscription Credentials :
Credentials required for access to the Azure Subscription.

The ID assigned by Azure for the Subscription that is monitored.
Subscription ID (default is: ): 09x73b6b-bcxb-40x3-92xd-ebx7-EXAMPLE

The tenant ID that is assigned by Azure. Used to log in to the Azure service API.
Tenant ID (default is: ): 75x2e745-e4x4-42x7-94xb-dex1-EXAMPLE

The client ID that is assigned by Azure to identify this agent as an external
application that monitors the Azure compute services.
Client ID (default is: ): 79x2e6c3-7exd-41x2-a9x9-4fx2-EXAMPLE

The secret access key or password that is created by Azure for the client application.
Enter Secret Key/Password (default is: ): hidden

Re-type : Secret Key/Password (default is: ): hidden

Azure Compute Virtual Machine Subnode :

Create agent subnodes to define groupings of virtual machines. Each subnode name
must be unique within your environment. It is suggested that you have no more
than 50 virtual machines per subnode.

No 'Azure Compute Virtual Machine Subnode' settings available.
Edit 'Azure Compute Virtual Machine Subnode' settings, [1=Add, 2=Edit, 3=Del,
4=Next, 5=Exit] (default is: 5): 1
Subnode Name (default is: ): azc1
The type of filter to be applied.
Filter Type [ 1=All, 2=Tag Name-Value Pair, 3=Resource Group ] (default is: 1): 2

The filter value corresponding to the selected Filter Type. This value can be a
Resource Group or Tag Name-Value Pair, for example Environment\Production.
A backslash might appear in the example, do not enter a backslash in the value
you provide.
Filter Value (default is: ): Environment:Production

Azure Compute Virtual Machine Subnode settings: Subnode Name=azc1
Edit 'Azure Compute Virtual Machine Subnode' settings, [1=Add, 2=Edit, 3=Del,
4=Next, 5=Exit] (default is: 5): 5
Configuration completed successfully.
Automatic start at system initialization has been configured.
Automatic stop at system shutdown has been configured.
You have new mail in /var/spool/mail/root
```

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

The silent response file contains the agent configuration parameters with default values that are defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters.

After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

Follow these steps to configure the Azure Compute agent in the silent mode.

1. Open the `azure_compute_silent_config.txt` file at one of the following paths in a text editor.

- **Linux** `install_dir/samples/azure_compute_silent_config.txt`  
Example, `/opt/ibm/apm/agent/samples/azure_compute_silent_config.txt`
- **Windows** `install_dir\samples\azure_compute_silent_config.txt`  
Example, `C:\IBM\APM\samples\azure_compute_silent_config.txt`

Where `install_dir` is the path where the agent is installed.

2. In the `azure_compute_silent_config.txt` file, specify values for all mandatory parameters and modify the default values of other parameters as needed.

See “[Configuration parameters for the Azure Compute agent](#)” on page 244 for an explanation of each of the configuration parameters.

**Important:** You must enable and specify the Filter Type and Filter Value parameters for at least one subnode name.

3. Save and close the `azure_compute_silent_config.txt` file, and run the following command:

- **Linux** `install_dir/bin/azure_compute-agent.sh config instance-name install_dir/samples/azure_compute_silent_config.txt`  
Example, `/opt/ibm/apm/agent/bin/azure_compute-agent.sh config azc-inst3 /opt/ibm/apm/agent/samples/azure_compute_silent_config.txt`
- **Windows** `install_dir\bin\azure_compute-agent.bat config instance-name install_dir\samples\azure_compute_silent_config.txt`  
Example, `C:\IBM\APM\bin\azure_compute-agent.bat config azc-inst3 C:\IBM\APM\samples\azure_compute_silent_config.txt`

Where `install_dir` is the path where the agent is installed and `instance-name` is the name that you want to give to the agent instance.

4. Run the following command to start the agent:

- **Linux** `install_dir/bin/azure_compute-agent.sh start instance-name`  
Example, `/opt/ibm/apm/agent/bin/azure_compute-agent.sh start azc-inst3`
- **Windows** `install_dir\bin\azure_compute-agent.bat start instance-name`  
Example, `C:\IBM\APM\bin\azure_compute-agent.bat start azc-inst3`

Where `install_dir` is the path where the agent is installed and `instance-name` is the name of the agent instance.

## Example

Edited `azure_compute_silent_config.txt` with three subnodes that are named `account-all`, `env-prod`, and `LG1`.

```
#
# This is a sample configuration response file for agent Azure Compute.
#
# It contains an entry for every configuration property.
# Entries for optional properties that have no default value are included in
# comments.
# Entries for subnode AVM are given a sample subnode instance name of avm1.
# Ensure that all uncommented properties have a value before configuring the
# agent.
#
# Subscription ID: The ID assigned by Azure for the Subscription that is
# monitored.
KAK_SUBSCRIPTION_ID=09x73b6b-bcxb-40x3-92xd-ebx7-EXAMPLE
# Tenant ID: The tenant ID that is assigned by Azure. Used to log in to the
# Azure service API.
KAK_TENANT_ID=75x2e745-e4x4-42x7-94xb-dex1-EXAMPLE
```

```

# Client ID: The client ID that is assigned by Azure to identify this agent
# as an external
# application that monitors the Azure compute services.
KAK_CLIENT_ID=79x2e6c3-7exd-41x2-a9x9-4fx2-EXAMPLE
# Secret Key/Password: The secret access key or password that is created by
# Azure for the client application.
KAK_SECRET_PASSWORD=hZxwPq/IOx1nvg/wdxLwTf2Fs3x2sWQV/sCE-EXAMPLE

# Filter Type: The type of filter to be applied.
# Valid values: ALL (All), TAG_NAME_VALUE (Tag Name-Value Pair),
# RESOURCE_GROUP (Resource Group)
#KAK_FILTER_TYPE.avm1=ALL
# Filter Value: The filter value corresponding to the selected Filter Type.
# This value can be a Resource Group or Tag Name-Value Pair, for example
# Environment:Production. A backslash might appear in the example, do not
# enter a backslash in the value you provide.
#KAK_FILTER_VALUE.avm1=

# Filter Type: The type of filter to be applied.
# Valid values: ALL (All), TAG_NAME_VALUE (Tag Name-Value Pair),
# RESOURCE_GROUP (Resource Group)
KAK_FILTER_TYPE.account-all=ALL
# Filter Value: The filter value corresponding to the selected Filter Type.
# This value can be a Resource Group or Tag Name-Value Pair, for example
# Environment:Production. A backslash might appear in the example, do not
# enter a backslash in the value you provide.
KAK_FILTER_VALUE.account-all=

# Filter Type: The type of filter to be applied.
# Valid values: ALL (All), TAG_NAME_VALUE (Tag Name-Value Pair),
# RESOURCE_GROUP (Resource Group)
KAK_FILTER_TYPE.env-prod=TAG_NAME_VALUE
# Filter Value: The filter value corresponding to the selected Filter Type.
# This value can be a Resource Group or Tag Name-Value Pair, for example
# Environment:Production. A backslash might appear in the example, do not
# enter a backslash in the value you provide.
KAK_FILTER_VALUE.env-prod=DTAP:prod

# Filter Type: The type of filter to be applied.
# Valid values: ALL (All), TAG_NAME_VALUE (Tag Name-Value Pair),
# RESOURCE_GROUP (Resource Group)
KAK_FILTER_TYPE.LG1=RESOURCE_GROUP
# Filter Value: The filter value corresponding to the selected Filter Type.
# This value can be a Resource Group or Tag Name-Value Pair, for example
# Environment:Production. A backslash might appear in the example, do not
# enter a backslash in the value you provide.
KAK_FILTER_VALUE.LG1=linux-group1

```

### What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring Cisco Unified Computing System (UCS) monitoring

You must configure the Cisco UCS agent to monitor the health, network, and performance of Cisco Unified Computing System (UCS).

### Before you begin

- Review the hardware and software prerequisites. For the up-to-date system requirement information, see the [“System requirements” on page 75](#) for Cisco UCS agent.
- Ensure that the user, who connects to the Cisco Unified Computing System Manager (UCSM) infrastructure, has administrator privileges. You can use an existing user ID, which has administrator privileges, or create a new user ID by completing the steps that are mentioned in the [“Creating a user and granting required permissions” on page 251](#) section.
- If the Cisco UCS agent is configured to communicate with its Cisco UCS data sources that use the SSL agent, add the SSL certificate of each data source to the certificate truststore of the agent. For more information about enabling SSL communication with Cisco UCS data sources, see [“Enabling SSL communication with Cisco UCS data sources” on page 255](#).

### About this task

The Cisco UCS agent is a multiple instance agent. You must create the first instance, and start the agent manually. The directions here are for the most current release of this agent. For more information about how to check the version of an agent in your environment, see [Agent version command](#).

The configuration attributes define which Cisco UCS infrastructure is monitored. The attributes define a connection to Cisco UCSM 1.4, or later. You can configure more than one instance of the monitoring agent on a remote monitoring host system. You can also create separate instances to monitor specific Cisco UCS infrastructure.

After the Cisco UCS agent is installed, you can start the agent. However, you must manually configure the agent to view data for all the agent attributes.

- To configure the agent on Windows systems, you can use the **IBM Performance Management** window.
- To configure the agent on Linux systems, you can run the script and respond to prompts, or use the silent response file.

## Creating a user and granting required permissions

Before you configure the Cisco UCS agent, you must create a user and grant required permissions to the user to monitor the Cisco Unified Computing Systems (UCS).

### Procedure

1. Open the **Red Hat Enterprise Virtualization Manager Web Administration** portal.
2. Click **Configure**.
3. In the **Configuration** window, select **Roles**.
  - a) To create a role, click **New**.
  - b) In the **New Role** window, add the name of the role and select **Admin** as the account type.
  - c) Ensure that the check boxes in the **Check boxes to Allow Action** pane are not selected, and click **OK**.
4. In the **Configuration** window, select **System Permission**.
  - a) To grant a user permission, click **Add**.
  - b) In the **Add System Permission to User** window, select the user to whom you want to grant the permission.
  - c) From the **Assign role to user** list, select the role that you created and click **OK**.

## Configuring the agent on Windows systems

You can configure the Cisco UCS agent on Windows operating systems by using the **IBM Performance Management** window. After you update the configuration values, you must start the agent to save the updated values.

### About this task

You can configure the agent when the agent is running or stopped. The agent remains in the same state after configuration. For example, if the agent is running, it remains in the running state after configuration. The Cisco UCS agent provides default values for some parameters. You can specify different values for these parameters.

### Procedure

To configure the agent on Windows systems, follow these steps:

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Monitoring Agent for Cisco UCS**, and then click **Configure agent**.

**Remember:** After you configure the agent for the first time, the **Configure agent** option is disabled. To configure the agent again, click **Reconfigure**.

3. In the Monitoring Agent for Cisco UCS window, follow these steps:

- a) Enter a unique name for the Cisco UCS agent instance, and click **OK**.
- b) On the **CONFIG** tab, specify values for the configuration parameters, and then click **Next**.
- c) On the **LOG\_CONFIG** tab, specify values for the configuration parameters, and then click **Next**.

For more information about the configuration parameters in each tab of the Monitoring Agent for Cisco UCS window, see the following topics:

- [“Configuration parameters for the agent” on page 254](#)
- [“Configuration parameters for the data provider” on page 255](#)

4. In the **IBM Performance Management** window, right-click **Monitoring Agent for Cisco UCS**, and then click **Start**.

### What to do next

- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).
- If you are monitoring a large Cisco UCS environment, you might need to increase the heap size for the Java™ data provider. For more information, see [“Increasing the Java heap size” on page 256](#).

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

The silent response file contains the agent configuration parameters with default values that are defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters.

After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

- To configure the Cisco UCS agent in the silent mode, follow these steps:
  - a) In a text editor, open the `cisco_ucs_silent_config.txt` file that is available at the following path:

– **Linux** `install_dir/samples/cisco_ucs_silent_config.txt`

For example, `/opt/ibm/apm/agent/samples/cisco_ucs_silent_config.txt`

– **Windows** `install_dir\samples\cisco_ucs_silent_config.txt`

For example, `C:\IBM\APM\samples\cisco_ucs_silent_config.txt`

- b) In the `cisco_ucs_silent_config.txt` file, specify values for all mandatory parameters. You can also modify the default values of other parameters.

For more information about the configuration parameters, see the following topics:

- [“Configuration parameters for the agent” on page 254](#)
- [“Configuration parameters for the data provider” on page 255](#)

- c) Save and close the `cisco_ucs_silent_config.txt` file, and run the following command:



– **Linux** `install_dir/bin/cisco_ucs-agent.sh config instance_name  
install_dir/samples/cisco_ucs_silent_config.txt`

For example, `/opt/ibm/apm/agent/bin/cisco_ucs-agent.sh config  
instance_name /opt/ibm/apm/agent/samples/cisco_ucs_silent_config.txt`

– **Windows** `install_dir\bin\cisco_ucs-agent.bat config instance_name  
install_dir\samples\cisco_ucs_silent_config.txt`

For example, `C:\IBM\APM\bin\cisco_ucs-agent.bat config instance_name C:\IBM  
\APM\samples\cisco_ucs_silent_config.txt`

Where,

***instance\_name***

Name that you want to give to the instance.

***install\_dir***

Path where the agent is installed.

**Important:** Ensure that you include the absolute path to the silent response file. Otherwise, the agent data is not shown in the dashboards.

d) Run the following command to start the agent:

– **Linux** `install_dir/bin/cisco_ucs-agent.sh start instance_name`

For example, `/opt/ibm/apm/agent/bin/cisco_ucs-agent.sh start instance_name`

– **Windows** `install_dir\bin\cisco_ucs-agent.bat start instance_name`

For example, `C:\IBM\APM\bin\cisco_ucs-agent.bat start instance_name`

Where,

***instance\_name***

Name that you want to give to the instance.

***install\_dir***

Path where the agent is installed.

## What to do next

- Log in to the Cloud APM console to view the data that is collected by the agent in the dashboards. For more information about using the Cloud APM console, see [“Starting the Cloud App Management UI” on page 176](#).
- If you are monitoring a large Cisco UCS environment, you might need to increase the heap size for the Java™ data provider. For more information, see [“Increasing the Java heap size” on page 256](#).

## Configuring the agent on Linux systems

To configure the agent on Linux operating systems, you must run the script and respond to prompts.

### Procedure

- To configure the agent by running the script and responding to prompts, follow these steps:

a) Go to command line and enter the following command:

`install_dir/bin/cisco_ucs-agent.sh config instance_name`

For example, `/opt/ibm/apm/agent/bin/cisco_ucs-agent.sh config instance_name`

Where,

***instance\_name***

Name that you want to give to the instance.

**install\_dir**

Path where the agent is installed.

b) Respond to the prompts by referring to the following topics:

- [“Configuration parameters for the agent” on page 254](#)
- [“Configuration parameters for the data provider” on page 255](#)

c) Run the following command to start the agent:

```
install_dir/bin/cisco_ucs-agent.sh start instance_name
```

For example, **/opt/ibm/apm/agent/bin/cisco\_ucs-agent.sh start instance\_name**

**What to do next**

- Log in to the Cloud APM console to view the data that is collected by the agent in the dashboards. For more information about using the Cloud APM console, see [“Starting the Cloud App Management UI” on page 176](#).
- If you are monitoring a large Cisco UCS environment, you might need to increase the heap size for the Java data provider. For more information, see [“Increasing the Java heap size” on page 256](#).

**Configuration parameters for the agent**

When you configure the Cisco UCS agent, you can change the default values of the configuration parameters, such as the instance name and the SSL validation certificates.

The following table contains detailed description of the configuration parameters for the Cisco UCS agent.

<i>Table 27. Name and description of the configuration parameters for the Cisco UCS agent</i>		
<b>Parameter name</b>	<b>Description</b>	<b>Mandatory field</b>
Instance Name	The name of the instance.  <b>Restriction:</b> The <b>Instance Name</b> field displays the name of the instance that you specify when you configure the agent for the first time. When you configure the agent again, you cannot change the instance name of the agent.	Yes
URL	The URL of the Cisco UCS Manager.	Yes
User name	The administrator user name of the Cisco UCS Manager.	Yes
Password	The administrator password of the Cisco UCS Manager.	Yes
Confirm Password	The same password that you entered in the <b>Password</b> field.	Yes
SSL truststore filepath	The path of the secure socket layer (SSL) truststore file.  If you want the agent to validate SSL certificates when you use SSL to communicate over the network, then specify the location where the secure socket layer (SSL) truststore file is located.	Yes

Table 27. Name and description of the configuration parameters for the Cisco UCS agent (continued)

Parameter name	Description	Mandatory field
Validate SSL Certificates	<p>A Boolean value that indicates whether the agent validates SSL certificates when the agent uses SSL to communicate over the network.</p> <p>Set the value to Yes if you want the agent to validate SSL certificates when the agent uses SSL to communicate over the network. Set the value to No to prevent the agent from validating SSL certificates.</p> <p><b>Tip:</b> For more information about enabling SSL communication with Cisco UCS data sources, see <a href="#">“Enabling SSL communication with Cisco UCS data sources”</a> on page 255.</p>	Yes

## Configuration parameters for the data provider

When you configure the Cisco UCS agent, you can change the default values of the parameters for the data provider, such as the maximum number of data provider log files, the maximum size of the log file, and the level of detail that is included in the log file.

The following table contains detailed descriptions of the configuration parameters for the data provider.

Table 28. Name and description of the configuration parameters for the data provider

Parameter name	Description	Mandatory field
Maximum number of Data Provider Log Files	The maximum number of log files that the data provider creates before it overwrites the previous log files. The default value is 10.	Yes
Maximum Size in KB of Each Data Provider Log	The maximum size in KB that a data provider log file must reach before the data provider creates a new log file. The default value is 5190 KB.	Yes
Level of Detail in Data Provider Log	The level of detail that can be included in the log file that the data provider creates. The default value is INFO. The following values are valid: OFF, SEVERE, WARNING, INFO, FINE, FINER, FINEST, and ALL.	Yes

## Enabling SSL communication with Cisco UCS data sources

The Cisco UCS agent can be configured to securely communicate with its Cisco Unified Computing System (UCS) data sources by using SSL. In this configuration, you must add a data source SSL certificate to the certificate truststore of the agent.

### About this task

**Important:** The following information applies only if the agent is configured to validate SSL certificates.

If SSL certificate validation is turned off, the Cisco UCS agent connects to Cisco UCS data sources even if the SSL certificates are expired, untrusted, or invalid. However, turning off SSL certificate validation is potentially not secure and must be done with care.

If a Cisco UCS data source uses an SSL certificate that is signed by a common certificate authority, then it is not necessary to add certificates to the agent certificate truststore. However, if the data source uses a certificate that is not signed by a common certificate authority, then add the certificate to the truststore. Doing so allows the agent to connect and collect data.

## Procedure

1. Copy the certificate file from your data source to the agent computer.
2. On the agent computer, place the certificate file in a directory of your choice. Do not overwrite the certificate files. Use a unique file name and label for each certificate that you add.
3. Use the `keytool` command to add the data source certificate to the certificate truststore of the agent:

```
keytool -import -noprompt -trustcacerts -alias CertificateAlias -file \  
CertificateFile -keystore Truststore -storepass TruststorePassword
```

Where,

### ***CertificateAlias***

Unique reference for each certificate added to the certificate truststore of the agent. For example, an appropriate alias for the certificate from *datasource.example.com* is *datasource*.

### ***CertificateFile***

Complete path and file name to the Cisco UCS data source certificate to add to the truststore.

### ***Truststore***

Complete path and file name to the Cisco UCS agent certificate database. Use the following path and file name:

- **Windows** (64 bit) *install\_dir\tmaitm6\_x64\kv6.truststore*
- **Linux** (64 bit) *install\_dir/lx8266/vm/etc/kv6.truststore*

### ***TruststorePassword***

ITMFORVE is the default password for the Cisco UCS agent truststore. To change the password, refer the Java Runtime documentation.

**Important:** To use the `keytool` command, the Java Runtime bin directory must be in your path. Use the following commands:

- **Windows** (64 bit) `set PATH=%PATH%;install_dir\java\java70_x64\jre\bin`
- **Linux** (64 bit) `PATH="$PATH":/opt/ibm/apm/agent/JRE/lx8266/bin`

4. After you add all the data source certificates, start the monitoring agent.

## Increasing the Java heap size

After you configure the Cisco UCS agent, if you are monitoring a large Cisco UCS environment, then you might need to increase the heap size for the Java™ data provider.

### **About this task**

The default heap size for the Java data provider is 256 megabytes. In large Cisco UCS environments, if the following problems arise, then you might need to increase the heap size:

- The Java data provider stops because of a `javacore` problem, and creates a file that is named `javacore.date.time.number.txt` in the `CANDLEHOME\tmaitm6_x64` directory.
- The `javacore.date.time.number.txt` file contains the string `java/lang/OutOfMemoryError`.

## Procedure

- **Windows**  
Complete the following steps to set a value of 1 GB as heap size:
  1. Open the `%CANDLE_HOME%\TMAITM6_x64\kv6_data_provider.bat` file.

2. Add the following line before the line that starts with `KV6_JVM_ARGS="$KV6_CUSTOM_JVM_ARGS . . .`:

```
SET KV6_CUSTOM_JVM_ARGS=-Xmx1024m
```

3. Restart the agent.

- **Linux**

Complete the following steps to set a value of 1 GB as heap size:

1. Open the `$CANDLEHOME/1x8266/vm/bin/kv6_data_provider.sh` file.
2. Add the following line before the line that starts with `KV6_JVM_ARGS="$KV6_CUSTOM_JVM_ARGS . . .`:

```
KV6_CUSTOM_JVM_ARGS=-Xmx1024m
```

3. Restart the agent.

## Configuring Cassandra monitoring

You must configure the agent after installation so that the agent can collect data from nodes within the cluster to monitor the health and performance of Cassandra database.

### Before you begin

Ensure that the system requirements for the Cassandra agent are met in your environment.

### About this task

The Cassandra agent is a multiple instance agent. You must create the first instance and start the agent instance manually. You can configure the agent on Windows and Linux operating systems.

To configure the agent on Windows systems, you can use the IBM Cloud Application Performance Management window, or the silent response file.

To configure the agent on Linux systems, you can run the command-line configuration script and respond to its prompts, or the silent response file.

The product version and the agent version often differ. The directions here are for the most current release of this agent. For more information about how to check the version of an agent in your environment, see [Agent version command](#).

## Configuring the agent on Windows systems

To configure the agent on Windows operating systems, you can use the IBM® Performance Management window. After you update the configuration values, start the agent to apply the updated values.

### Procedure

To configure the agent on Windows operating systems, complete the following steps:

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Template** in the **Task/SubSystem** column, and click **Configure Using Defaults**.

The **Monitoring Agent for Cassandra** window opens.

3. In the **Enter a unique instance name** field, type an agent instance name and click **OK**.
4. In the **Monitoring Agent for Cassandra** window, specify values for the configuration parameters and click **OK**.

For more information about the configuration parameters, see [“Configuration parameters of the agent” on page 259](#).

5. In the **IBM Performance Management** window, right-click the agent instance and click **Start**.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the console, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent on Linux systems

To configure the agent on Linux operating systems, you can run the command line configuration script and respond to its prompts.

### Procedure

To configure the agent on Linux operating systems, complete the following steps:

1. On the command line, change the path to the agent installation directory.  
For example, `/opt/ibm/apm/agent/bin`
2. Run the following command, where instance name is the name that you want to give to the instance:  
`./cassandra-agent.sh config instance_name`
3. When the command line displays the following message, type 1 and enter:  
Edit 'Monitoring Agent for Cassandra' setting? [1=Yes, 2=No]
4. Specify values for the configuration parameters when you are prompted.  
For more information about the configuration parameters, see [“Configuration parameters of the agent” on page 259](#).
5. Run the following command to start the agent:  
`./cassandra-agent.sh start instance_name`

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the console, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

You can use the silent response file to configure the Cassandra agent on Linux and Windows systems. After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

To configure the agent in the silent mode, complete the following steps:

1. In a text editor, open the silent configuration file that is available at the following location and specify values for all the parameters:

**Windows** `install_dir\samples\cassandra_silent_config_windows.txt`

**Linux** `install_dir\samples\cassandra_silent_config_UNIX.txt`

**Windows** `C:\IBM\APM\samples`

**Linux** `/opt/ibm/apm/agent/samples`

For more information about the configuration parameters, see [“Configuration parameters of the agent” on page 259](#).

2. On the command line, change the path to `install_dir\bin`.

3. Run the following command:

```
Windows cassandra-agent.bat config instance_name install_dir\samples  
\cassandra_silent_config_windows.txt
```

```
Linux cassandra-agent.sh config instance_name install_dir\samples  
\cassandra_silent_config_UNIX.txt
```

4. Start the agent.

```
Windows In the IBM Performance Management window, right-click the agent instance that you  
created, and click Start.
```

```
Linux Run the following command: ./cassandra-agent.sh start instance_name
```

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the console, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuration parameters of the agent

When you configure the Cassandra agent, you can change the default value of the parameters, such as IP address and JMX\_PORT.

The following table contains detailed description of the configuration parameters of the Cassandra agent.

Parameter name	Description	Mandatory field
Instance Name	The default value for this field is identical to the value that you specify in the <b>Enter a unique instance name</b> field.	Yes
IP Address	The IP address of any node in the cluster.	Yes
JMX_PORT	The Java Management Extension (JMX) port number to enable monitoring. <b>Important:</b> Ensure that you specify the JMX port, JMX username, and JMX password throughout the cluster. If the node through which the agent is connected to the cluster is not working, then the agent collects data from a different node in the cluster by using the same parameters.	Yes
JMX_Username	The username for accessing the JMX.	No
JMX_Password	The password for accessing the JMX.	No

## Configuring Citrix Virtual Desktop Infrastructure monitoring

The Monitoring Agent for Citrix Virtual Desktop Infrastructure monitors the following functions: Citrix XenDesktop component, Event log and alerts, and Citrix XenDesktop services. Additionally, you can view the Load Index Summary metrics performance data for Citrix XenApp and XenDesktop. You can diagnose problematic login times by viewing the performance data for the login steps.

### Before you begin

- Make sure that the system requirements for the Citrix VDI agent are met in your environment. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\) for the Citrix VDI agent](#).
- Ensure that the following information is available:
  - Host name of the delivery controller to which you plan to connect.

- OData user name, password, and domain.
- PowerShell user name, password, domain, PowerShell port, SSL verification type, and authentication mechanism if you enable Windows Event Log Event and PowerShell metric retrieval.
- Ensure that an agent operator user account has at least Citrix read-only administrator privileges. See [Enabling Citrix read-only administrator privileges](#).
- Starting with Citrix VDI agent version 8.1.3.1, the ability to retrieve Windows Event Log Events became available. To retrieve Windows Event Log Events from all Desktop Delivery Controller (DDC) and Virtual Delivery Agent (VDA) machines, remote PowerShell access needs to be enabled for the user account that is specified during the agent instance configuration. Follow these steps to ensure that the agent can perform this function:
  1. Log in to a Windows computer as the user specified in the agent instance configuration.
  2. Run the following PowerShell command, where *vda\_system* is the name of a VDA machine that is powered on:
 

```
Get-WinEvent -FilterHashtable
@{ProviderName='Citrix*';LogName='Citrix*';StartTime=((Get-Date).AddDays(-10))} -ComputerName vda_system
```
- Ensure that the following load balancing policies are enabled for the monitored environment:
  - CPU Usage
  - Disk Usage
  - Memory Usage

These policies can be configured through the Citrix Studio application.

### About this task

The Citrix VDI agent is a multiple instance agent. You must create at least one instance, and start the agent instance manually.

The configuration for XenApp servers is the same as for XenDesktop servers. If a configuration parameter name or description mentions only "XenDesktop", it is also for XenApp.

### Procedure

1. Configure the agent on Windows systems with the **IBM Performance Management** window or the silent response file.
  - [“Configuring the agent on Windows systems”](#) on page 262.
  - [“Configuring the agent by using the silent response file”](#) on page 264.
2. Configure the agent on Linux systems with the script that prompts for responses or the silent response file.
  - [“Configuring the agent by responding to prompts”](#) on page 263.
  - [“Configuring the agent by using the silent response file”](#) on page 264.

### What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI”](#) on page 176.

If you are unable to view the data in the agent dashboards, first check the server connection logs and then the data provider logs. The default paths to these logs are listed here:

- **Linux** /opt/ibm/apm/agent/logs
- **Windows** C:\IBM\APM\TMAITM6\_x64\logs



## Configuration parameters for the Citrix VDI agent

The configuration parameters for the Citrix VDI agent are displayed in a table.

### 1. [Citrix VDI Agent Settings](#) - Citrix VDI agent environment settings.

Parameter name	Description	Silent configuration file parameter name
XenApp or XenDesktop Site Name	Provide a name to identify the XenApp or XenDesktop site agent instance. Example, <i>vdi_inst2</i>  <b>Note:</b> This alias can be anything that you choose to represent the WebLogic server agent instance with the following restrictions. Only letters, Arabic numerals, the underline character, and the minus character can be used in the connection name. The maximum length of a connection name is 25 characters.	Each of the following parameters must have an instance name suffix that is the same for each parameter of an agent instance. New agent instances must use a unique instance name for its set of parameters. For example, one agent instance can use <i>.vdi1</i> and another agent instance can use <i>.vdi2</i> in place of <i>.instance_name</i> in the parameter names that follow.
Delivery Controller	Host name or IP address of the delivery controller. If multiple DDCs are set up in a cluster, a ' ' separated list of delivery controllers can be provided.	<b>KVD_XDS_DELIVERY_CONTROLLER.instance_name</b>
User Name	User name that is used to authenticate with the OData API on the specified XenApp or XenDesktop delivery controller.	<b>KVD_XDS_ODATA_USERNAME.instance_name</b>
Password	Password that is used to authenticate with the OData API on the specified XenApp or XenDesktop delivery controller.	<b>KVD_XDS_ODATA_PASSWORD.instance_name</b>
Domain	Domain that is used to authenticate with the OData API on the specified XenApp or XenDesktop delivery controller.	<b>KVD_XDS_ODATA_DOMAIN.instance_name</b>
PowerShell User name	User name that is used to authenticate for PowerShell calls to remote VDA and DDC machines.  <b>Note:</b> This and all following PowerShell parameters are only needed when <a href="#">“Enabling monitoring of Windows events and PowerShell metrics”</a> on page 266. These advanced environment variables are off by default because of the significant load they put on the monitored system.	<b>KVD_XDS_POWERSHELL_USERNAME.instance_name</b>
PowerShell Password	Password that is associated with PowerShell user name provided.	<b>KVD_XDS_POWERSHELL_PASSWORD.instance_name</b>
PowerShell Domain	Domain that is associated with PowerShell user name provided.	<b>KVD_XDS_POWERSHELL_DOMAIN.instance_name</b>

Table 30. Citrix VDI Agent Settings (continued)

Parameter name	Description	Silent configuration file parameter name
PowerShell Port	The SSL port that is open for use by WinRm. PowerShell default remote connection ports are 5985 for HTTP and 5986 for HTTPS.	<b>KVD_XDS_POWERSHELL_PORT.instance_name</b>
SSL Requirement	Choose the SSL option required for your environment.	<b>KVD_XDS_SSL_CONFIG.instance_name</b> Valid values, <b>KVD_XDS_SSL_CONFIG_VERIFY</b> Verify <b>KVD_XDS_SSL_CONFIG_NOVERIFY</b> No Verify <b>KVD_XDS_SSL_CONFIG_NOSSL</b> No SSL
PowerShell Authentication Mechanism	Defines the type of authentication that is used to create a credential to retrieve information from remote systems with PowerShell.	<b>KVD_XDS_POWERSHELL_AUTH_MECH.instance_name</b> Valid values, <b>KVD_XDS_POWERSHELL_BASIC</b> Basic <b>KVD_XDS_POWERSHELL_CREDSSP</b> CredSSP <b>KVD_XDS_POWERSHELL_NTLM</b> NTLM <b>KVD_XDS_POWERSHELL_DEFAULT</b> Default <b>KVD_XDS_POWERSHELL_DIGEST</b> Digest <b>KVD_XDS_POWERSHELL_KERBEROS</b> Kerberos <b>KVD_XDS_POWERSHELL_NEGOTIATE</b> Negotiate

## Configuring the agent on Windows systems

You can configure the Citrix VDI agent on Windows operating systems by using the IBM Cloud App Management window. After you update the configuration values, you must start the agent to save the updated values.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Cloud App Management**.
2. In the **IBM Performance Management** window, right-click the **Monitoring Agent for Citrix Virtual Desktop Infrastructure** template, and then click **Configure agent**.

**Remember:** After you configure an agent instance for the first time, the **Configure agent** option is disabled. To configure the agent instance again, right-click on it and then click **Reconfigure...**

3. In the Monitoring Agent for Citrix Virtual Desktop Infrastructure window, complete the following steps:

- a) Enter a unique instance name for the Monitoring Agent for Citrix Virtual Desktop Infrastructure instance, and click **OK**.
4. Click **Next** on the agent instance name window.
5. Enter the **XenApp and XenDesktop Site Configuration** instance template settings.
 

**Note:** This section is not the XenApp or XenDesktop site instance configuration. It is a template section for setting what is used as the default values when you add the actual XenApp or XenDesktop site instance configurations in [step 6](#).

See [Table 30 on page 261](#) for an explanation of each of the configuration parameters.
6. Press **New** and enter XenApp or XenDesktop site instance settings, then click **Next**.
 

See [Table 30 on page 261](#) for an explanation of each of the configuration parameters.

**Note:** The **PowerShell User name** parameter and all following PowerShell parameters are only needed when [“Enabling monitoring of Windows events and PowerShell metrics” on page 266](#). These advanced environment variables are off by default because of the significant load they put on the monitored system.

**Note:** Ensure the **SSL Config** and **PowerShell Authentication Mechanism** parameters are set correctly for each new XenApp or XenDesktop site instance. A defect causes the default values to be set instead of the template values.
7. Click **OK** to complete the configuration.
8. In the IBM Cloud App Management window, right-click the instance that you configured, and then click **Start**.

## Configuring the agent by responding to prompts

After installation of the Citrix VDI agent, you must configure it before you start the agent. If the Citrix VDI agent is installed on a local Linux machine, you can follow these instructions to configure it interactively through command line prompts.

### About this task

**Remember:** If you are reconfiguring a configured agent instance, the value that is set in the last configuration is displayed for each setting. If you want to clear an existing value, press the space key when the setting is displayed.

### Procedure

Follow these steps to configure the Citrix VDI agent by running a script and responding to prompts.

1. Run the following command:

```
install_dir/bin/citrixvdi-agent.sh config instance_name
```

where *install\_dir* is the path where the agent is installed and *instance\_name* is the name that you want to give to the agent instance.

Example

```
/opt/ibm/apm/agent/bin/citrixvdi-agent.sh config vdi_inst01
```

2. Respond to the prompts to set configuration values for the agent.
 

See [“Configuration parameters for the Citrix VDI agent” on page 261](#) for an explanation of each of the configuration parameters.

**Note:** The **PowerShell User name** parameter and all following PowerShell parameters are only needed when [“Enabling monitoring of Windows events and PowerShell metrics” on page 266](#). These advanced environment variables are off by default because of the significant load they put on the monitored system.
3. Run the following command to start the agent:

```
install_dir/bin/citrixvdi-agent.sh start instance_name
```

where *install\_dir* is the path where the agent is installed and *instance\_name* is the name of the agent instance.

Example

```
/opt/ibm/apm/agent/bin/citrixvdi-agent.sh start vdi_inst01
```

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

The silent response file contains the agent configuration parameters with default values that are defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters.

After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

- Configure the Citrix VDI agent in the silent mode:
  - a) Open the `citrixvdi_silent_config.txt` file at one of the following paths in a text editor.

- **Linux** `install_dir/samples/citrixvdi_silent_config.txt`  
Example, `/opt/ibm/apm/agent/samples/citrixvdi_silent_config.txt`
- **Windows** `install_dir\samples\citrixvdi_silent_config.txt`  
Example, `C:\IBM\APM\samples\citrixvdi_silent_config.txt`

where *install\_dir* is the path where the agent is installed.

- b) In the `citrixvdi_silent_config.txt` file, specify values for all mandatory parameters and modify the default values of other parameters as needed.

See [“Configuration parameters for the Citrix VDI agent” on page 261](#) for an explanation of each of the configuration parameters.

**Note:** The **PowerShell User name** parameter and all following PowerShell parameters are only needed when [“Enabling monitoring of Windows events and PowerShell metrics” on page 266](#). These advanced environment variables are off by default because of the significant load they put on the monitored system.

- c) Save and close the `citrixvdi_silent_config.txt` file, and run the following command:

- **Linux** `install_dir/bin/citrixvdi-agent.sh config instance_name install_dir/samples/citrixvdi_silent_config.txt`  
Example, `/opt/ibm/apm/agent/bin/citrixvdi-agent.sh config vdi_inst01 /opt/ibm/apm/agent/samples/citrixvdi_silent_config.txt`
- **Windows** `install_dir\bin\citrixvdi-agent.bat config instance_name install_dir\samples\citrixvdi_silent_config.txt`

Example, `C:\IBM\APM\bin\citrixvdi-agent.bat config vdi_inst01 C:\IBM\APM\samples\citrixvdi_silent_config.txt`

where *install\_dir* is the path where the agent is installed and *instance\_name* is the name that you want to give to the agent instance.

**Important:** Ensure that you include the absolute path to the silent response file. Otherwise, the agent data is not shown in the dashboards.

d) Run the following command to start the agent:

```
- Linux install_dir/bin/citrixvdi-agent.sh start instance_name  
Example, /opt/ibm/apm/agent/bin/citrixvdi-agent.sh start vdi_inst01
```

```
- Windows install_dir\bin\citrixvdi-agent.bat start instance_name  
Example, C:\IBM\APM\bin\citrixvdi-agent.bat start vdi_inst01
```

where *install\_dir* is the path where the agent is installed and *instance\_name* is the name of the agent instance.

## Enabling Citrix read-only administrator privileges

The Citrix VDI agent requires the agent operator user account have at least Citrix read-only administrator privileges.

### About this task

To run these steps remotely from a computer that has the Citrix Delegated Admin PowerShell Snap-in installed, use the `AdminAddress` parameter. For example, the command in step 2 would become `New-AdminAdministrator -Name "YOURDOMAIN\NewAdmin" -AdminAddress "controller1.YOURDOMAIN.com"`. Where *YOURDOMAIN* is the name of the network domain, *NewAdmin* is the user account that is being given Citrix administration privileges, and *controller1.YOURDOMAIN.com* is the fully qualified domain name of the Citrix site server.

### Procedure

1. Start a PowerShell session with an existing Citrix administrator account.
2. Load the Delegated Admin PowerShell Snap-in to manage the Citrix XenApp or XenDesktop site.

```
(Add-PSSnapin Citrix.DelegatedAdmin.Admin.V1)
```

3. Add the agent operator user account as a Citrix site administrator.

```
New-AdminAdministrator -Name "YOURDOMAIN\NewAdmin"
```

Where *YOURDOMAIN* is the name of the network domain and *NewAdmin* is the user account that is being given Citrix administration privileges.

4. Query for the available roles and scopes to assign to *NewAdmin*.

```
Get-AdminRole  
Get-AdminScope
```

5. Assign roles to the agent operator user account, including read-only administrator permissions.

```
Add-AdminRight -Administrator "YOURDOMAIN\NewAdmin" -Role "Read Only Administrator" -Scope "All"
```

Where

- *YOURDOMAIN* is the name of the network domain.
- *NewAdmin* is the user account that is being given Citrix administration privileges.
- *Read Only Administrator* is the Citrix site administrator role that you are assigning.
- *All* is the Citrix site administrator scope that you are assigning.

6. Confirm the addition and changes of the administrator.

```
Get-AdminAdministrator -Name "YOURDOMAIN\NewAdmin"
```

Where *YOURDOMAIN* is the name of the network domain and *NewAdmin* is the user account that is being given Citrix administration privileges.

## Enabling monitoring of Windows events and PowerShell metrics

Enable monitoring of Windows events and PowerShell metrics with this procedure. Monitoring this data can have a significant performance impact to the monitored system.

### Before you begin

Ensure the agent's [PowerShell configuration parameters](#) are set.

### About this task

One or more of the following advanced environment variables must be enabled for the agent to monitor Windows events and PowerShell metrics.

#### GET\_SESSION\_LATENCY

Whether session latency and round-trip time are retrieved remotely from the connected VDA from PowerShell.

#### GET\_VDA\_MACHINE\_METRICS\_REMOTELY

Whether VDA machine metrics are retrieved remotely from PowerShell.

#### RETRIEVE\_WINDOWS\_EVENTS

Whether Windows Event Log Events are retrieved from PowerShell from Windows VDAs and DDCs.

### Procedure

1. Go to the agent installation directory of the Citrix VDI agent:

- **Linux** `install_dir/config`
- **Windows** `install_dir\TMAITM6_x64`

where `install_dir` is the path where the agent is installed.

2. Edit the Citrix VDI agent configuration file to set one or more of the `GET_SESSION_LATENCY`, `GET_VDA_MACHINE_METRICS_REMOTELY`, and `RETRIEVE_WINDOWS_EVENTS` variables to `true`.

- **Linux** `vd.environment`
- **Windows** `KVDENV_instance_name`

where `instance_name` is the name of the agent instance.

3. Restart the agent.

**Important:** To make these settings the default for all new agent instances, set them to `true` in the configuration template files:

- **Linux** This setting is already made the default for new agents instances by editing `vd.environment` in [Step 2](#).
- **Windows** `KVDENV`

### Example

```
GET_SESSION_LATENCY=true
GET_VDA_MACHINE_METRICS_REMOTELY=true
RETRIEVE_WINDOWS_EVENTS=true
```

## 2019.4.0.2 Configuring CouchDB monitoring

You must configure the CouchDB agent so that the agent can collect data to monitor the availability and performance of CouchDB server resources.

### Before you begin

- Ensure that the system requirements for the CouchDB agent are met in your environment.

- Ensure that a user is created in CouchDB database to run the agent. The user must have an admin access on the CouchDB database that it monitors.
- Check the version of CouchDB database that you are using. The Monitoring agent for CouchDB supports CouchDB version 2.3.1 and above.

### About this task

The CouchDB agent is a multi instance agent. You must configure the agent manually after it is installed. You can configure the agent on Windows and Linux operating systems. The agent requires an instance name and the CouchDB server user credentials to configure it. The instance name that you specify can contain up to 28 characters.

## 2019.4.0.2 Configuring the agent on Windows systems

You can configure the agent on Windows operating systems by using the IBM Cloud Application Performance Management window. After you update the configuration values, start the agent to apply the updated values.

### Procedure

To configure the agent on Windows operating systems, complete the following steps:

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, complete these steps:
  - a) Double-click the **Monitoring Agent for CouchDB** template.
  - b) In the **Monitoring Agent for CouchDB** window, specify an instance name and click **OK**.
3. In the **Monitoring Agent for CouchDB** window, complete these steps:
  - a) In the **IP Address** field, enter the IP address of CouchDB server that you want to monitor.
  - b) In the **Port number** field, enter the port number on which CouchDB server is running. The default value is 5984.
  - c) In the **User name** field, enter the user name of CouchDB user.
  - d) In the **Password** field, enter the password of CouchDB user.
  - e) In the **Confirm Password** field, enter the password again.
4. Click **Next**.
5. In the **Java Parameters** window, enter following inputs:
  - a) From the **Java trace level** list, select a trace level for Java.  
The default value is Error.
  - b) Click **OK**.  
The instance is displayed in the **IBM Performance Management** window.
6. Right-click the **Monitoring Agent for CouchDB** instance, and click **Start**.

**Remember:** To configure the agent again, complete either of the two steps given below in the **IBM Performance Management**

Reconfigure the existing agent instance in **IBM Performance Management**:

- a. Stop the agent instance that you want to configure.
- b. Right-click the **Monitoring Agent for CouchDB** instance, and click **Reconfigure**.
- c. Repeat steps 3 to 5.

Configure by creating a new agent instance in **IBM Performance Management**:

- a. Repeat steps 2 to 5.

## What to do next

- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

### 2019.4.0.2 Configuring the agent on Linux systems

You can run the configuration script and respond to prompts to configure the agent on Linux operating systems.

#### Procedure

To configure the agent on Linux operating systems, complete the following steps:

1. On command line, run the following command:

```
install_dir/bin/couchdb-agent.sh config instance_name
```

Where *instance\_name* is the name you want to give to the instance, and *install\_dir* is the installation directory for the CouchDB agent.

2. When you are prompted to enter a value for the following parameters, press **Enter** to accept the default value or specify a different value and press **Enter**.
  - IP address
  - Port number (Default value is 5984)
  - User name
  - Password
  - Re-type Password
  - Java trace level (Default value is `ERROR`)
3. Run the following command to start the agent:

```
install_dir/bin/couchdb-agent.sh start instance_name
```

## What to do next

- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

### 2019.4.0.2 Configuring the agent by using the silent response file

Use the silent response file to configure the agent without responding to prompts when you run the configuration script. You can use the silent response file for configuring the agent on both Windows and Linux systems.

#### About this task

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the silent configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode. You can use the silent response file for configuring the agent on both Windows and Linux systems.

#### Procedure

1. In a text editor, open the response file that is available at the following path:

```
Linux install_dir/samples/couchdb_silent_config.txt
```



**Windows** `install_dir\samples\couchdb_silent_config.txt`

Where `install_dir` is the installation directory of the CouchDB agent.

2. In the response file, specify a value for the following parameters:

- For the **IP Address** parameter, specify the IP address of a CouchDB server that you want to monitor.
- For the **Port number** parameter, specify a port number. The default value is 5984.
- For the **User name** parameter, enter the user name.
- For the **Password** parameter, enter the user password.
- For the **Re-type Password** parameter, enter the user password again.
- For the **Java trace level** parameter, retain the default value as `ERROR`.

3. Save and close the response file, and run the following command to update the agent configuration settings:

**Linux** `install_dir/bin/couchdb-agent.sh config instance_name install_dir/samples/couchdb_silent_config.txt`

**Windows** `install_dir\BIN\couchdb-agent.bat config instance_name install_dir\samples\couchdb_silent_config.txt`

Where `instance_name` is the name that you want to give to the instance, and `install_dir` is the installation directory of CouchDB.

**Important:** Be sure to include the absolute path to the silent response file. Otherwise, no agent data is displayed in the dashboards.

4. Start the agent using the following command:

**Linux** | **UNIX** Run the following command: `install_dir\bin\couchdb-agent.sh start`

**Windows** Right-click **Monitoring Agent for CouchDB** and then click **Start**.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## 2019.4.0.2 Configuration parameters for the agent

When you configure the CouchDB agent, you can change the default values of the parameters, such as the instance name and the SSL validation certificates.

The following table contains detailed descriptions of the configuration parameters for the CouchDB agent.

Parameter name	Description	Mandatory field
IP Address	The IP address of the node where the CouchDB is installed.	Yes
Username	The user name of the CouchDB user.	Yes
Password	The password to connect to the CouchDB server.	Yes
Confirm Password	The same password that you entered in the <b>Password</b> field.	Yes
Port Number	The port number where CouchDB is listening. Use the default port number 5984, or specify another port number.	Yes

Parameter name	Description	Mandatory field
Java trace level	The trace level of the Java provider. The valid trace level values are as follows: <ul style="list-style-type: none"> <li>• OFF</li> <li>• ERROR</li> <li>• WARN</li> <li>• INFO</li> <li>• DEBUG_MAX</li> <li>• ALL</li> </ul>	Yes

## Configuring WebLogic monitoring

The Monitoring Agent for WebLogic provides you with a central point of monitoring for the health, availability, and performance of your WebLogic server environment. The agent displays a comprehensive set of metrics to help you make informed decisions about your WebLogic resources, including Java virtual machines (JVMs), Java messaging service (JMS), Java Database Connectivity (JDBC).

### Before you begin

- The product version and the agent version often differ. The directions here are for the most current release of this agent. To access the documentation for earlier agent releases, see the [Table 32 on page 270](#) table.

Agent version	Documentation
8.1.4, 8.1.4.1	<a href="#">8.1.4</a>
8.1.3.2	<a href="#">IBM Performance Management 8.1.3</a> <sup>1</sup>

<sup>1</sup> The link opens an on-premises Knowledge Center topic.

- Make sure that the system requirements for the WebLogic agent are met in your environment. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\) for the WebLogic agent](#).
- Before you configure the WebLogic agent, the Oracle WebLogic server first must be configured by completing the following tasks:

**Note:** Most of the Oracle WebLogic server configuration is done by using the administrative console, typically at <http://weblogic-server:7001/console>.

1. Set up a monitor user in the Monitors group.
  - a. Select the domain to monitor/edit.
  - b. Select **Security Realms**.
  - c. Select your security realm (or create one if one does not exist).
  - d. Create a user that will be used to communicate with WebLogic over JMX.
  - e. Add this user the Monitors group.
  - f. Save the changes to the domain.
2. Enable the Listen Ports.
  - a. Select the domain to monitor/edit.

- b. On each server that you want to monitor, click **Environment > Servers > Select a server**.
  - c. Ensure that the **Listen Port** is enabled and note its port number.
  - d. If you want to enable SSL, then ensure that the **SSL Listener Port** is enabled and set a port for SSL as well.
3. Enable the JMX MBean Server Connections.
    - a. Select the domain that you want to monitor/edit.
    - b. Select **Configure > Advanced**.
    - c. Check **Platform Mbean Server Enabled**.
    - d. Save the change.
  4. Enable the IIOP Protocol option.
    - a. Select the domain that you want to monitor/edit.
    - b. On each server that you would like to monitor, click **Environment > Servers** then select a server.
    - c. Select the **Protocol Tab > Select IIOP**.
    - d. Under the **Advanced** section, enter the default IIOP user name and password.
    - e. Save the change.
  5. Enable SSL.
    - a. Enable HTTP Tunneling.
      - 1) Go to **Environment > Servers > Select a server > Protocol > General**.
      - 2) Check **Enable HTTP Tunneling**.
    - b. Enable SSL Listen Port.
      - 1) Go to **Environment > Servers > Select a server > Configuration > General**.
      - 2) Configure a port number.

### About this task

The WebLogic agent is both a multiple instance agent and also a multiple subnode agent. You can create one agent instance with multiple subnodes – one for each WebLogic server, or you can create an agent instance for each WebLogic server with one subnode for that server. Or you can create a combination of each type of configuration. After you configure agent instances, you must start each agent instance manually.

### Procedure

1. To configure the agent on Windows systems, use the **IBM Performance Management** window or the silent response file with the agent configuration batch file.
  - [“Configuring the agent on Windows systems” on page 272.](#)
  - [“Configuring the agent by using the silent response file” on page 276.](#)
2. To configure the agent on Linux and UNIX systems, run the agent configuration script and respond to prompts, or use the silent response file.
  - [“Configuring the agent by responding to prompts” on page 276.](#)
  - [“Configuring the agent by using the silent response file” on page 276.](#)

### What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176.](#)

If you are unable to view the data in the agent dashboards, first check the server connection logs and then the data provider logs. The default paths to these logs are as follows.

- **Linux** /opt/ibm/apm/agent/logs
- **Windows** C:\IBM\APM\TMAITM6\_x64\logs

## Configuration parameters for the WebLogic agent

The configuration parameters for the WebLogic agent are displayed in a table.

1. WebLogic Agent Settings - WebLogic agent environment settings.

Parameter name	Description	Silent configuration file parameter name
WebLogic Server Name	Provide a name to identify the WebLogic server agent instance. Example: <i>wls1</i>  <b>Note:</b> This alias can be anything you choose to represent the WebLogic server agent instance with the following restrictions. Only letters, Arabic numerals, the underline character, and the minus character can be used in the connection name. The maximum length of a connection name is 25 characters.	Each of the following parameters must have an instance name suffix which will be the same for each parameter of an agent instance. New agent instances must use a unique instance name for its set of parameters. For example, one agent instance can use <i>.wls1</i> and another agent instance can use <i>.wls2</i> in place of <i>.instance_name</i> in the parameter names below.
User Name	Username that is used to authenticate with the WebLogic server.	<b>KWB_WLS_USERNAME.instance_name</b>
Password	Password that is used to authenticate with the WebLogic server.	<b>KWB_WLS_PASSWORD.instance_name</b>
Host	Host that is used to authenticate with the WebLogic server. Type either the fully qualified host name or the IP address of the WebLogic server.	<b>KWB_WLS_HOST.instance_name</b>
Port	Port that is used to authenticate with the WebLogic server.	<b>KWB_WLS_PORT.instance_name</b>
Protocol	Protocol that is used to authenticate with the WebLogic server. Supported protocols are <i>iiop</i> and <i>https</i> .	<b>KWB_WLS_PROTOCOL.instance_name</b>

## Configuring the agent on Windows systems

You can configure the WebLogic agent on Windows operating systems by using the **IBM Performance Management** window. After you update the configuration values, you must start the agent to save the updated values.

### Procedure

1. Click **Start > All Programs > IBM Cloud App Management > .**
2. In the **IBM Performance Management** window, right-click the **Monitoring Agent for WebLogic** template, and then click **Configure agent**.

**Remember:** After you configure an agent instance for the first time, the **Configure agent** option is disabled. To configure the agent instance again, right-click on it and then click **Reconfigure....**

3. Enter a unique instance name then click **OK**. Use only letters, Arabic numerals, the underline character, and the minus character in the instance name. For example: *weblogic01*.

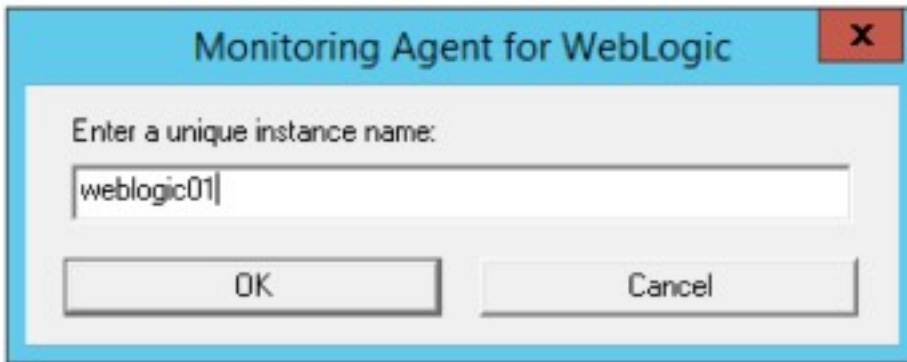


Figure 1. The window to enter a unique instance name

4. Click **Next** on the **Instance Name** agent configuration panel.

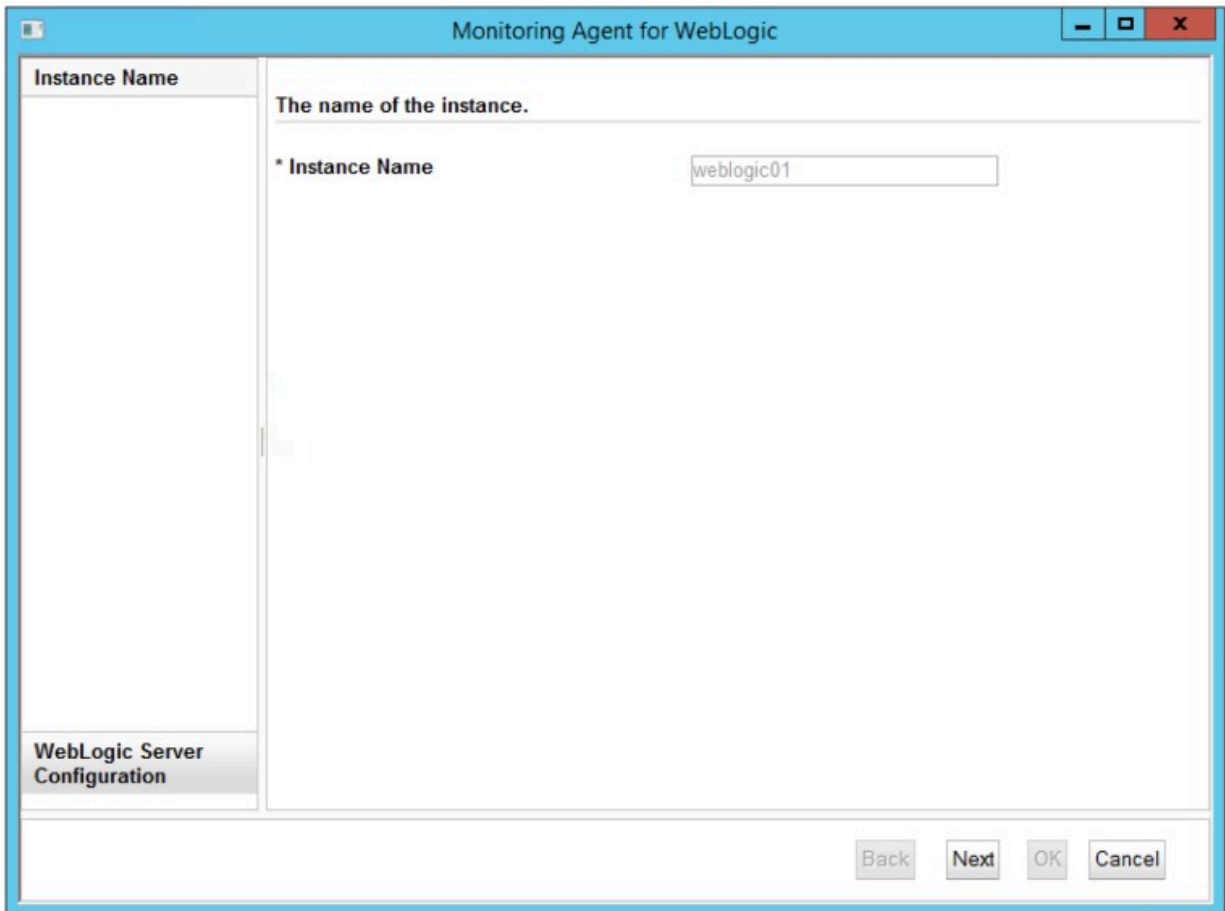


Figure 2. The window displaying the agent instance name

5. Enter the **WebLogic Server Configuration** instance template settings.

**Note:** This section is not the WebLogic server connection instance configuration. It is a template section for setting what is used as the default values when you add the actual WebLogic server connection instance configurations beginning in [step 6](#).

See [Table 33](#) on page 272 for an explanation of each of the configuration parameters.

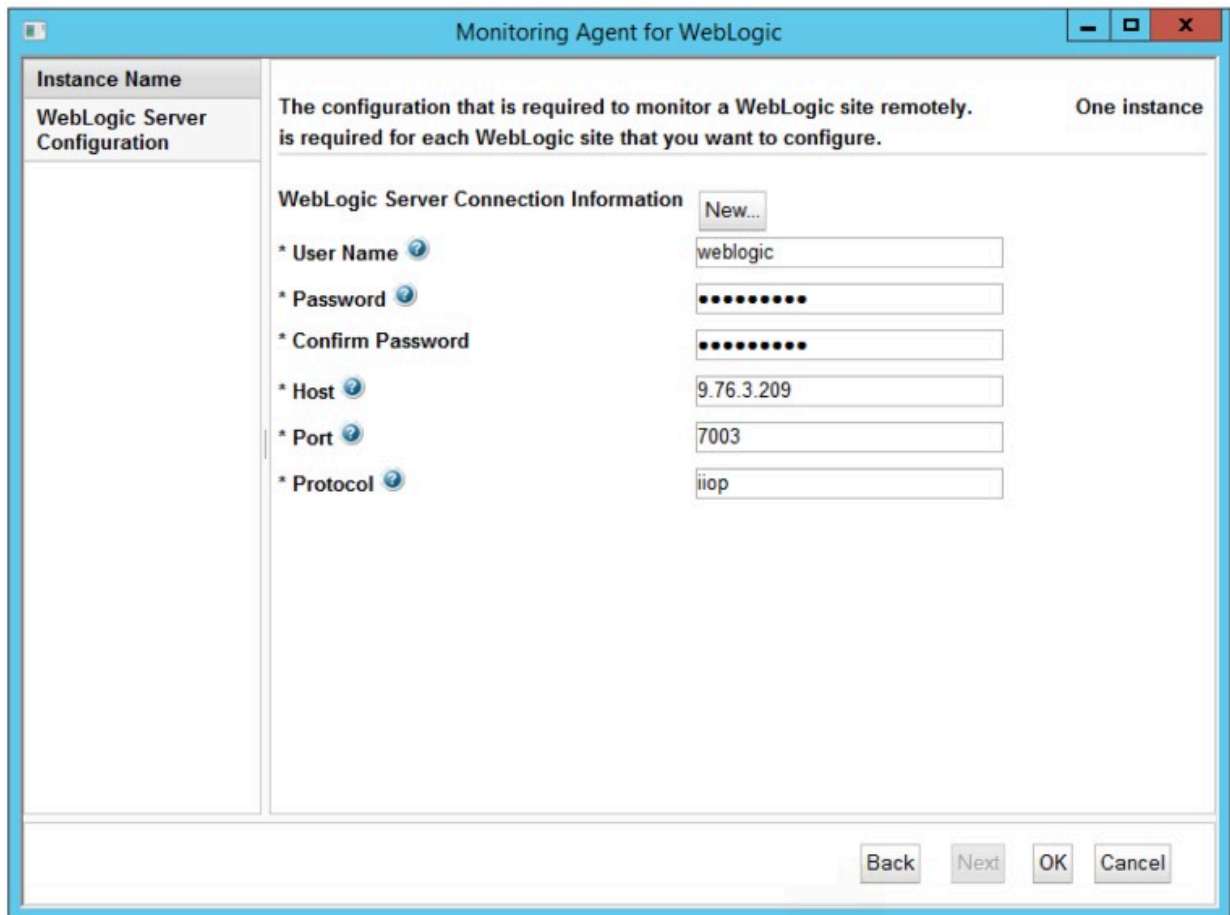


Figure 3. The window to specify WebLogic server connection instance template settings

6. Press **New** and enter WebLogic server connection instance settings, then click **Next**. See [Table 33](#) on page 272 for an explanation of each of the configuration parameters.

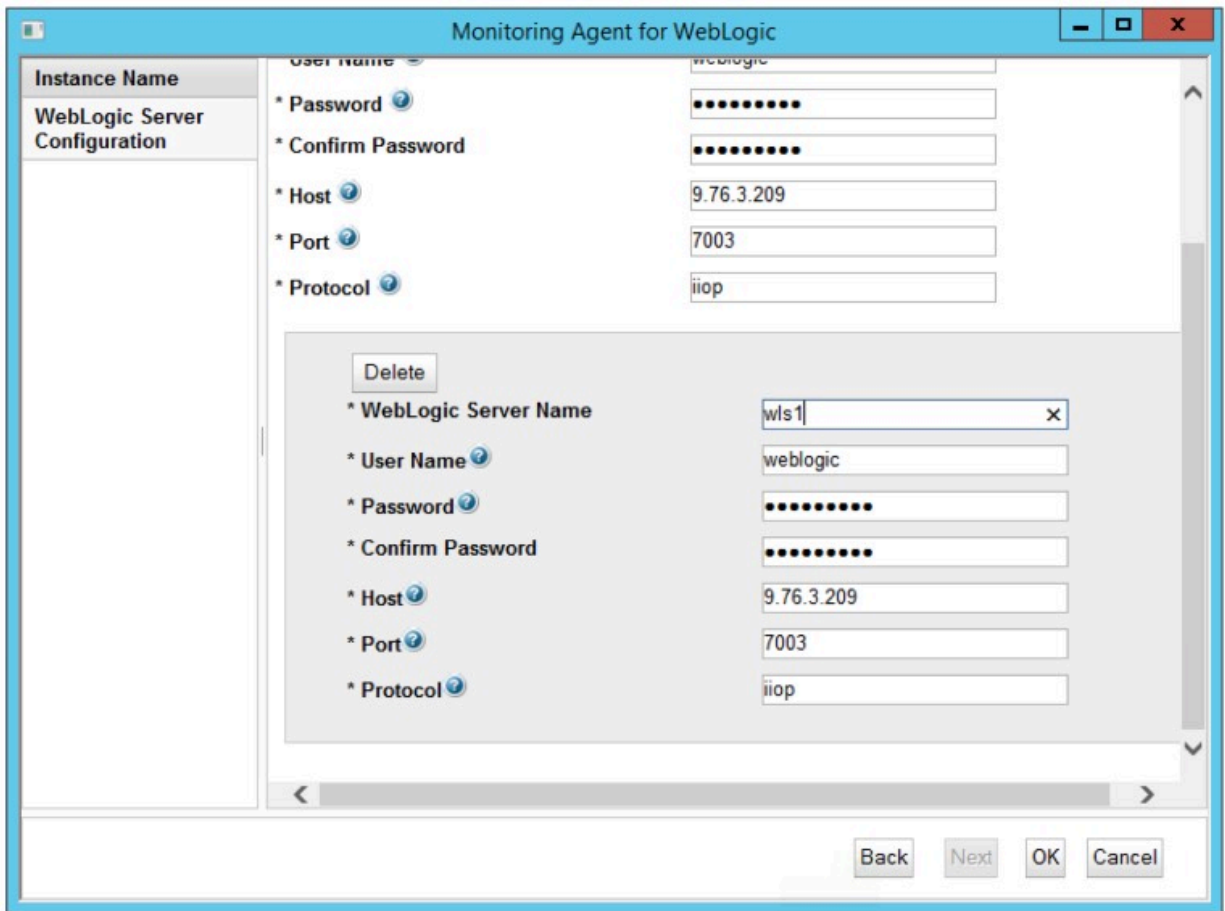


Figure 4. The window to specify WebLogic server connection instance settings

7. Click **OK** to complete the configuration.
8. Copy the WebLogic security files into the WebLogic agent binary directory.
  - a. Locate the `wlclient.jar` and `wljmxclient.jar` files under `ORACLE_HOME`. For example, `C:\Oracle\Middleware\Oracle_Home\wlserver\server\lib`.
  - b. Copy the files from step “8.a” on page 275 to the WebLogic agent binary directory.
    - **Linux** | **UNIX** `install_dir/bin`.
    - **Windows** `install_dir\TMAITM6_x64`
 where `install_dir` is the path where the agent is installed. The default `install_dir` paths are listed here:
    - **Linux** | **UNIX** `/opt/ibm/apm/agent`
    - **Windows** `C:\IBM\APM\TMAITM6_x64`
9. In the IBM Cloud App Management window, right-click the instance that you configured, and then click **Start**.

## Configuring the agent by responding to prompts

After installation of the WebLogic agent, you must configure it before you start the agent. If the WebLogic agent is installed on a local Linux or UNIX computer, you can follow these instructions to configure it interactively through command line prompts.

### About this task

**Remember:** If you are reconfiguring a configured agent instance, the value that is set in the last configuration is displayed for each setting. If you want to clear an existing value, press the space key when the setting is displayed.

### Procedure

Follow these steps to configure the WebLogic agent by running a script and responding to prompts.

1. Run the following command.

```
install_dir/bin/weblogic-agent.sh config instance_name
```

where *install\_dir* is the path where the agent is installed and *instance\_name* is the name that you want to give to the agent instance.

Example

```
/opt/ibm/apm/agent/bin/weblogic-agent.sh config example-inst01
```

2. Respond to the prompts to set configuration values for the agent.

See [“Configuration parameters for the WebLogic agent”](#) on page 272 for an explanation of each of the configuration parameters.

3. Copy the WebLogic client library files into the WebLogic agent binary directory.

- a) Locate the `wlclient.jar` and `wljmxclient.jar` files under `ORACLE_HOME`.
- b) Copy the files from step [“3.a”](#) on page 276 to the WebLogic agent binary directory.

```
install_dir/bin
```

where *install\_dir* is the path where the agent is installed.

Example

```
/opt/ibm/apm/agent/bin
```

4. Run the following command to start the agent:

```
install_dir/bin/weblogic-agent.sh start instance_name
```

where *install\_dir* is the path where the agent is installed and *instance\_name* is the name of the agent instance.

Example

```
/opt/ibm/apm/agent/bin/weblogic-agent.sh start example-inst01
```

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

The silent response file contains the agent configuration parameters with default values defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters.



After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

## Procedure

Configure the WebLogic agent in the silent mode by completing the following steps.

1. In a text editor, open the `weblogic_silent_config.txt` file that is available at the following path:

- **Linux** | **UNIX** `install_dir/samples/weblogic_silent_config.txt`
- **Windows** `install_dir\samples\weblogic_silent_config.txt`

where `install_dir` is the path where the agent is installed.

Example

- **Linux** | **UNIX** `/opt/ibm/apm/agent/samples/weblogic_silent_config.txt`
- **Windows** `C:\IBM\APM\samples\weblogic_silent_config.txt`

2. In the `weblogic_silent_config.txt` file, specify values for all mandatory parameters. You can also modify the default values of other parameters.

See “[Configuration parameters for the WebLogic agent](#)” on page 272 for an explanation of each of the configuration parameters.

3. Save and close the `weblogic_silent_config.txt` file, and run the following command:

- **Linux** | **UNIX** `install_dir/bin/weblogic-agent.sh config instance_name install_dir/samples/weblogic_silent_config.txt`
- **Windows** `install_dir\bin\weblogic-agent.bat config instance_name install_dir\samples\weblogic_silent_config.txt`

where `install_dir` is the path where the agent is installed and `instance_name` is the name that you want to give to the agent instance.

The default `install_dir` paths are listed here:

- **Linux** | **UNIX** `/opt/ibm/apm/agent`
- **Windows** `C:\IBM\APM\TMAITM6_x64`

**Important:** Ensure that you include the absolute path to the silent response file. Otherwise, the agent data is not shown in the dashboards.

Example

- **Linux** | **UNIX** `/opt/ibm/apm/agent/bin/weblogic-agent.sh config example-inst01 /opt/ibm/apm/agent/samples/weblogic_silent_config.txt`
- **Windows** `C:\IBM\APM\bin\weblogic-agent.bat config example-inst01 C:\IBM\APM\samples\weblogic_silent_config.txt`

4. Copy the WebLogic client libraries into the WebLogic agent binary directory.

- a. Locate the `wlclient.jar` and `wljmxclient.jar` files under `ORACLE_HOME`.
- b. Copy the files from step “5.a” on page 277 to the WebLogic agent binary directory.

- **Linux** | **UNIX** `install_dir/bin.`
- **Windows** `install_dir\TMAITM6_x64`

where `install_dir` is the path where the agent is installed. The default `install_dir` paths are listed here:

- **Linux** | **UNIX** `/opt/ibm/apm/agent`
- **Windows** `C:\IBM\APM\TMAITM6_x64`

5. Run the following command to start the agent:

- **Linux** | **UNIX** `install_dir/bin/weblogic-agent.sh start instance_name`
- **Windows** `install_dir\bin\weblogic-agent.bat start instance_name`

where *install\_dir* is the path where the agent is installed and *instance\_name* is the name of the agent instance.

The default *install\_dir* paths are listed here:

- **Linux** | **UNIX** `/opt/ibm/apm/agent`
- **Windows** `C:\IBM\APM\TMAITM6_x64`

Example

- **Linux** | **UNIX** `/opt/ibm/apm/agent/bin/weblogic-agent.sh start example-inst01`
- **Windows** `C:\IBM\APM\bin\weblogic-agent.bat start example-inst01`

## Configuring DataPower monitoring

To monitor DataPower appliances, you need to first complete some configuration tasks on your appliances, and then configure the Monitoring Agent for DataPower.

The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history” on page 52](#).

## Configuring DataPower appliances

Before you configure the Monitoring Agent for DataPower, you must complete some configuration tasks on your appliances.

**Tip:** For information about the supported DataPower appliances, see the Prerequisites tab in [Software Product Compatibility Reports](#) [\[2\]](#).

To see monitoring data, such as resource utilization, throughput, and connection statistics, enable resource monitoring. For instructions, see [“Resource monitoring” on page 278](#).

**Important:** Make sure that the user ID has the proper permissions to configure the DataPower appliance. You can enter `*/*/*?Access=x` in the **Access profile** field for the user ID that is used to configure the DataPower appliance. And then use this user ID to configure the DataPower appliance.

### Exporting the public certificate

If the XML Management Interface of the DataPower appliance has the SSL Proxy Profile enabled, you must export the public certificate to the machine that runs the DataPower agent.

#### Procedure

1. To download the crypto certificate, for example, `pubcert:///mycert.pem`, which is used by the XML Management Interface of the DataPower appliance, click **Administration > Main > File Management** and save the certificate to the machine that runs the DataPower agent.
2. When you configure the DataPower agent, an option to specify the **SSL Proxy Profile** field is available. Enter the absolute path of the public certificate.

### Resource monitoring

The first level of monitoring available for a DataPower appliance is to enable resource monitoring, such as SOAP management, statistics, and transaction rates.

The operation on the DataPower Gateway user interface (UI) in the following configuration tasks apply to DataPower Gateway Version 7.5.1 and former versions. If the version of the DataPower Gateway that you use is later than V 7.5.1, you can click the question mark icon in the UI and choose **WebGUI** to return to

the UI of the former version. And then follow the instructions to complete DataPower appliance configuration tasks.

### **Enabling SOAP management**

If you want the DataPower agent to collect data from DataPower Appliances, you must configure the XML Management Interface and enable SOAP Management.

#### **Procedure**

To enable SOAP:

1. Log on to the WebGUI for the DataPower Appliance that you want to monitor.
2. Click **Objects > Device Management > XML Management Interface**.

**Note:** Ensure that the Administrative state is enabled.

3. For **Port Number**, enter the port number on which the DataPower agent listens for notification reports. The port number is 5550 by default.
4. For **Enabled Services**, ensure that **SOAP Management** is selected.

### **Enabling Statistics**

If you want the DataPower agent to collect data from DataPower appliances, Statistics must be enabled.

#### **Procedure**

To enable Statistics, complete the following steps:

1. Log on to the WebGUI for the DataPower appliance that you want to monitor.
2. Click **Administration > Device > Statistics Settings**.
3. Enable **Statistics Settings** and click **Apply**.

### **Enabling Transaction Rate**

If you want the DataPower agent to collect data from DataPower appliances, the Transaction Rate must be enabled.

#### **Procedure**

To enable Transaction Rate, complete the following steps:

1. Log on to the WebGUI for the DataPower appliance that you want to monitor.
2. Select the default domain.
3. Click **Status > Connection > Transaction Rate**.
4. If **Statistics is currently disabled** is displayed, click **disabled** and in the Statistic Settings, set the **Administrative state to enabled**.
5. If you have multiple domains, click **Show All Domains** and repeat steps 3-4 to enable the Transaction Rate for all applicable domains.
6. Click **Apply**.

## **Configuring the DataPower agent**

The Monitoring Agent for DataPower provides a central point of monitoring for the DataPower appliances in your enterprise environment. You can identify and receive notifications about common problems with the appliances. The agent also provides information about performance, resource, and workload for the appliances.

### **About this task**

The DataPower agent is a multiple instance agent; you must create the first instance and start the agent manually. The Managed System Name includes the instance name that you specify, for example, *instance\_name: host\_name: pc*, where *pc* is your two character product code. The Managed System Name is limited to 32 characters.

The instance name that you specify is limited to 28 characters, minus the length of your host name. For example, if you specify DataPower as your instance name, your managed system name is DataPower:hostname:BN.

**Important:** If you specify a long instance name, the Managed System name is truncated and the agent code does not display correctly.

For each production DataPower appliance, configure one instance. If your DataPower appliances are non-production or small ones, you can configure only one agent instance to monitor them all. Multiple instances can run on the same machine. You can run the configuration script to create an instance and change any configuration settings. You can edit the agent silent response file before you run the script to bypass the prompts and responses that are required.

## Procedure

- To configure the DataPower agent, complete one of the following procedures:
  - **Linux** | **UNIX** To configure the agent by responding to prompts, complete the following steps:
    1. Go to the *install\_dir*/bin directory, where *install\_dir* is the installation directory for the DataPower agent.
    2. Run the `./datapower-agent.sh config instance_name` command.  
Choose an *instance\_name* that is unique on the server.
    3. When prompted to edit the DataPower agent settings, enter 1 to proceed.
    4. When prompted to edit the **Managed System Details**, enter one of the following options:
      - 1=Add
      - 2=Edit
      - 3=Del
      - 4=Next
      - 5=Exit

If it is the first time that you configure a DataPower agent instance on your system, the No 'DataPower Appliances' settings available message is displayed. Enter 1 to add a DataPower appliance setting. The default is option 5=Exit.

5. Enter the properties for the DataPower appliance:

### Managed System Name

For **Managed System Name**, enter the managed system name of the agent.

Choose a **Managed System Name** that is unique among all instances of the agent and that can be used to easily identify an appliance. The name should contain only alphanumeric characters, for example, the host name of the DataPower appliance.

### Device Host

For **Device Host**, enter the IP address of the monitored DataPower appliance.

### XML Management Interface Port

For **XML Management Interface Port**, enter the port number for the XML Management Interface. The default number is 5550.

### User ID

For **User ID**, enter the User ID to log in to the monitored DataPower appliance. The default value is admin.

### Password

For **Password**, enter the password to log in to the monitored DataPower appliance and then confirm the password.

### SSL Proxy Profile

For **SSL Proxy Profile**, enter the absolute path of the public certificate for your SSL proxy profile, if the XML management interface of the device is configured to use the profile. For example,

```
the location of the .pem file exported from datapower appliances/mycert.pem
```

where *the location of the .pem file exported from datapower appliances* is the absolute path of the public certificate. To export the public certificate, see [Exporting public certificate](#).

### SSL Proxy Option

For **SSL Proxy Option**, set to Yes if the XML management interface of the monitored device is configured to use a custom SSL proxy profile. Otherwise, set it to No.

6. To monitor multiple DataPower appliances, repeat “4” on page 280 and “5” on page 280 to configure one agent instance for each DataPower appliance. Otherwise, type 5 and press **Enter** to complete the configuration.
7. Run the following command to start the agent:

```
./datapower-agent.sh start instance_name
```

- Silent configuration

1. To configure the agent by editing the silent response file and running the script with no interaction, complete the following steps:

- **Linux** **UNIX** Open `install_dir/samples/datapower_silent_config.txt` in a text editor.

2. To configure the DataPower agent to monitor an appliance, enter the following properties:

#### Device Host

Enter the host name or IP address of the device. For example,  
**SOAP\_HOST.ManageSystemName= datapower01.**

#### XML Management Interface Port

Enter the port number for the XML Management Interface. The default value is 5550. For example, **DP\_PORT.ManageSystemName= 5550.**

#### User ID

Enter the User ID that is used to connect to the device. The default value is admin. For example, **DP\_UID.ManageSystemName= admin.**

#### Password

Enter the password of the User ID. For example, **DP\_PASSWORD.ManageSystemName= password.**

### SSL Proxy Profile

Enter the absolute path of the public certificate for your SSL proxy profile, if the XML management interface of the device is configured to use the profile. For example,

```
the location of the .pem file exported from datapower appliances/mycert.pem
```

where *the location of the .pem file exported from datapower appliances* is the absolute path of the public certificate. To export the public certificate, see [Exporting public certificate](#).

### SSL Proxy Option

For **SSL Proxy Option**, set to Yes if the XML management interface of the monitored device is configured to use a custom SSL proxy profile. Otherwise, set it to No. For example, **DP\_SSL\_OPTION.ManageSystemName1= Yes.**

**Important:** ManageSystemName is unique. You must replace it with your own system name in all entries. If you want to monitor multiple appliances, copy and repeat the steps that are shown to monitor an appliance. Remember to set the appropriate ManageSystemName and DataPower appliance parameters.

3. Go to the installation directory for the agent and run the following command to start the agent:

```
./datapower-agent.sh start instance_name
```

### What to do next

- To check the names and settings of the configured agent instances, run the `./cinfo -s bn` command.
- To display resource monitoring, configure the DataPower appliance accordingly. For instructions, see [Resource monitoring](#) of DataPower appliances.
- You can verify that the DataPower agent data is displayed in the UI.

## Configuring Db2 monitoring

The Monitoring Agent for Db2 monitors the availability and performance of the Db2 server. You can monitor multiple servers from the Cloud App Management; each server is monitored by a Db2 instance. Remote monitoring is also supported by Db2.

### Before you begin

Review hardware and software prerequisites. For the up-to-date system requirement information, see the [Software Product Compatibility Report](#).

**Note:** For Db2 agent execute `secure.sh` with root user only.

### About this task

The Db2 agent is a multiple instance agent, you must first create the instance and then start the agent manually.

The managed system name includes the agent instance name that you specify, for example, `instance_name:host_name:pc`.

Where:

- The `pc` is your two character product code.
- The `instance_name` is the agent instance name, and it must be the same as the Db2 instance name that is to be monitored.

The managed system name can contain up to 32 characters. The `instance_name` can contain up to 8 characters, excluding the length of your host name. For example, if you specify `DB2inst1` as your agent instance name, your managed system name is `DB2inst1:hostname:ud`.

Db2 agent has two resources DB2 Instances and DB2 Databases. The format for resource name for DB2 Instances is `<instance name>:<managed system name>`. This format enlists the DB2 instances. For example, if you specify `db2inst1` as your agent instance name, your DB2 Instance resource name is `db2inst1:db2inst1:hostname:ud`. The resource name for DB2 Databases has format as `<managed system name>: <instance name> : <database name>` which enlists all the databases present in DB2 Instances resource. For example, if you specify `db2inst1` as your agent instance name which has sample database, your DB2 Database resource name is `db2inst1:hostname:ud:db2inst1:SAMPLE`.

**Important:** If you specify a long agent instance name, the managed system name is truncated and the complete agent code is not displayed .

To avoid permission issues when you configure the agent, be sure to use the same root user or non-root user ID that was used for installing the agent. If you installed your agent as a selected user and want to configure the agent as a different user, see [“Starting agents as a non-root user”](#) on page 230.

Run the configuration script to create an instance and change the configuration settings. You can edit the Db2 silent response file before you run the configuration script to bypass the prompts and responses that are otherwise necessary.

After you configure the Db2 agent, be sure to start the agent with a user ID that has the Db2 SYSADM authority for the monitored instance. The agent requires the SYSADM authority to turn on all monitor

switches and collect the monitoring data. Therefore, a user with the SYSADM authority must start the agent. Use the instance owner user, which has the SYSADM authority, to start the agent.

## Procedure

To configure the agent with the default settings, complete the following steps:

1. Run the following command where *instance\_name* is the agent instance name:

```
install_dir/bin/db2-agent.sh config instance_name  
install_dir/samples/db2_silent_config.txt
```

**Note:** The agent instance name *instance\_name* is always the same as the Db2 instance name that is being monitored.

For more details about the existing agent instances, refer [“Viewing your managed resources” on page 769](#).

2. Run the following command to start the Db2 agent:

```
install_dir/bin/db2-agent.sh start instance_name
```

## What to do next

- Grant privileges to the Db2 user to view data for some attributes of the Db2. For information about granting these privileges, see [“Granting privileges for viewing Db2 agent metrics” on page 287](#).
- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent on Windows systems

You can use the IBM Cloud App Management window to configure the agent on Windows systems.

### Before you begin

Before you start configuring the Db2 agent for local and remote monitoring, ensure that following task is completed for remote monitoring.

- Set up client/server environment for remote monitoring, refer [“Prerequisites for Remote Monitoring” on page 291](#).

## Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Monitoring Agent for DB2**, and then click **Configure agent**.
3. In the **Enter a unique instance name** field, type the agent instance name and click **OK**.

**Important:** For local monitoring, the agent instance name must match the name of the Db2 instance that is being monitored.

For remote monitoring, the agent instance name must be the unique catalog node name.

4. In the **Monitoring Agent for DB2** window, complete these steps:

- a) In the **Username** field, enter the user name of Db2 instance.

For Local Db2, enter the name of Db2 instance owner.

For Remote Db2, enter Actual Db2 instance owner name from remote Db2 machine.

**Important:** This parameter is mandatory for monitoring remote Db2 instance.

- b) In the **Password** field, enter the password of Db2 instance.

For Local Db2, enter the password of Db2 instance owner.

For Remote Db2, enter Actual Db2 instance owner password from remote Db2 machine.

**Important:** This parameter is mandatory for monitoring remote Db2 instance.

- c) In the **DB2Customized SQL Definition File** field, enter the full file path name for the SQL definition file. If the SQL definition file is in the default directory, leave this field blank. Otherwise, enter the full file path name of the file. The default file name with path is as follows:

CANDLEHOME\TMAITM6\_x64\kudcussql.properties

- d) In the **db2diag Log File Path** field, enter the directory path for the db2diag log file. If the db2diag log file is in the default directory, leave this field blank. Otherwise, enter the path of the directory. The default directory path is as follows:

CANDLEHOME\TMAITM6\_x64\kudcussql.properties

**Note:** This parameter is not applicable for remote monitoring.

- e) In the **MSGID Filter in Regular Expression** field, enter the *MSGID* to filter the diagnostic log. The *MSGID* is a combination of the message type, message number, and severity level. Use a regular expression to filter the log based on message type, message number, or severity level, for example, ADM1\d\*1E|ADM222\d2W.
- f) From the **Enable Monitoring for Partitions in Remote Hosts** list, select Yes to specify that the Db2 agent can monitor partitions in remote hosts.
- g) From the **Enable Monitoring All Databases** list, select Yes to specify that the Db2 agent can monitor all databases.
- h) Click **OK**.

The agent instance is displayed in the IBM Cloud App Management window.

5. Run following steps to configure remote monitoring.

a) Open *install\_dir*\TMAITM6\_x64\KUDENV\_<instanceName>.

b) Set *KUD\_DB2\_CLIENT\_INST* to Db2 client instance name under which remote Db2 server instance is cataloged.

6. Right-click the **Monitoring Agent for DB2** instance, and click **Start**.

### What to do next

- Grant privileges to the Db2 user to view data for some attributes of the Db2 agent. For information about granting these privileges, see [“Granting privileges for viewing Db2 agent metrics” on page 287](#).
- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent on Linux or UNIX systems

Run the configuration script to configure the agent on Linux systems.

### Before you begin

Before you start configuring the Db2 agent for local and remote monitoring, ensure that following task is completed for remote monitoring.

- Set up client/server environment for remote monitoring, refer [“Prerequisites for Remote Monitoring” on page 291](#).

### Procedure

1. Run the following command

```
install_dir/bin/db2-agent.sh config instance_name
```

Where *instance\_name* is the name that you want to give to the instance:



**Important:** For local monitoring, the agent instance name must match the name of the Db2 instance that is being monitored.

For remote monitoring, the agent instance name must match the name of the local cataloged node of remote Db2 server instance that is to be monitored.

2. When you are prompted to provide a value for the following parameters, press Enter to accept the default value, or specify a value and then press Enter:

a) In the Username field, enter the user name of Db2 instance.

For Local Db2, enter the name of Db2 instance owner.

For Remote Db2, enter Actual Db2 instance owner name from remote Db2 machine.

**Important:** This parameter is mandatory for monitoring remote Db2 instance.

b) In the Password field, enter the password of Db2 instance.

For Local Db2, enter the password of Db2 instance owner.

For Remote Db2, enter Actual Db2 instance owner password from remote Db2 machine.

**Important:** This parameter is mandatory for monitoring remote Db2 instance.

c) In the DB2® SQL path field, enter the full file path name for the SQL definition file. If the SQL definition file is in the default directory, leave this field blank. Otherwise, enter the full file path name of the file. The default file name with path is as follows:

CANDLEHOME/config/kudcussql.properties

d) In the Diaglog path field, enter the directory path for the db2diag log file. If the db2diag log file is in the default directory, leave this field blank. Otherwise, enter the path of the directory. The default directory path is as follows:

/home/DB2owner\_home\_dir/sql1lib/db2dump

**Note:** This parameter is not applicable for remote monitoring.

e) In the Diaglog message ID filter field, enter the MSGID to filter the diagnostic log. The MSGID is a combination of the message type, message number, and severity level. Use a regular expression to filter the log based on message type, message number, or severity level, for example, ADM1\d\*1E|ADM222\d2W.

f) From the Monitor remote partitions list, enter Yes to specify that the Db2 agent can monitor partitions in remote hosts.

g) From the Monitor all databases list, enter Yes to specify that the Db2 agent can monitor all databases.

3. Run the following command to start the agent:

For local monitoring run

```
install_dir/bin/db2-agent.sh start instance_name
```

by Db2 instance owner user.

For remote monitoring, run

```
install_dir/bin/db2-agent.sh start node_name
```

with the instance owner of Db2 client instance under which remote Db2 server instance is cataloged.

### What to do next

- Grant privileges to the Db2 user to view data for some attributes of the Db2 agent. For information about granting these privileges, see [“Granting privileges for viewing Db2 agent metrics” on page 287](#).
- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent by using the silent response file

Use the silent response file to configure the agent without responding to prompts when you run the configuration script. You can use the silent response file for configuring the agent on both Windows and Linux systems.

### Before you begin

Before you start configuring the Db2 agent for local and remote monitoring, ensure that following task is completed for remote monitoring.

- Set up client/server environment for remote monitoring, refer [“Prerequisites for Remote Monitoring” on page 291](#).

### About this task

The silent response file contains the configuration parameters. You edit the parameter values in the response file, and run the configuration script to create an agent instance and update the configuration values.

### Procedure

1. In a text editor, open the `db2_silent_config.txt` file that is available at the following path:

```
Linux | UNIX | install_dir/samples/db2_silent_config.txt
```

```
Windows | install_dir\tmaitm6_x64\samples\db2_silent_config.txt
```

2. In the response file, specify a value for the following parameters:

- In the **Username**, enter the user name of Db2 instance.

For Local Db2, enter the name of Db2 instance owner.

For Remote Db2, enter Actual Db2 instance owner name from remote Db2 machine.

**Important:** This parameter is mandatory for monitoring remote Db2 instance.

- In the **Password**, enter the password of Db2 instance.

For Local Db2, enter the password of Db2 instance owner.

For Remote Db2, enter Actual Db2 instance owner password from remote Db2 machine.

**Important:** This parameter is mandatory for monitoring remote Db2 instance.

- For the **DB2 SQL path** parameter, leave this field blank if the SQL definition file is available at the default directory. Otherwise, enter the correct directory path. The SQL definition file is available at the following default path:

```
Linux | UNIX | CANDLEHOME/config/kudcussql.properties
```

For example, **KUD\_DB2\_SQL\_PATH=** `/opt/ibm/apm/agent/config/kudcussql.properties`

```
Windows | CANDLEHOME\TMAITM6_x64\kudcussql.properties
```

For example, **KUD\_DB2\_SQL\_PATH=** `C:\IBM\ITM\TMAITM6_x64\kudcussql.properties`

- For the **dialog path** parameter, leave this field blank if the `db2diag` log file is available at the default directory. Otherwise, enter the correct directory path. The log file is available at the following default path:

```
Linux | UNIX | /home/DB2owner_home_dir/sqlllib/db2dump
```

For example, **KUD\_DIAGLOG\_PATH=** `/home/db2inst1/sqlllib/db2dump`.

```
Windows | Windows Install_Driver:\ProgramData\IBM\DB2\DB2COPY  
\DB2INSTANCENAME
```

For example, **KUD\_DIAGLOG\_PATH=** `C:\ProgramData\IBM\DB2\DB2COPY1\DB2`

**Note:** This parameter is not applicable for remote monitoring.

- For the **dialog message ID filter** parameter, specify the *MSGID* to filter the diagnostic log. The MSGID is a combination of the message type, message number, and severity level. You can also use a regular expression, for example, **KUD\_DIAGLOG\_MSGID\_FILTER=** ADM1\d\*1E|ADM222\d2W.
  - For the **monitor remote partitions** parameter, enter Yes to specify that the Db2 agent monitors partitions in remote hosts. For example, **KUD\_MONITOR\_REMOTE\_PARTITIONS=** Yes.
  - For the **monitor all databases** parameter, enter Yes to specify that you want the Db2 agent to monitor all databases. For example, **KUD\_MONITOR\_ALL\_DATABASES=** Yes.
3. Save and close the `db2_silent_config.txt` file, and run the following command

```
Linux | UNIX | install_dir/bin/db2-agent.sh config instance_name
install_dir/samples/db2_silent_config.txt
Windows | install_dir\bin\db2-agent.bat config instance_name
\tmaitm6_x64\samples\db2_silent_config.txt
```

<instance\_name> is

- For monitoring Local Db2 server : The Db2 server instance name that you want to monitor.
- For monitoring Remote Db2 server: The catalog node name of remote Db2 server instance.

**Important:** Ensure that you include the absolute path to the silent response file. Otherwise, agent data is not shown in the dashboards.

4. For Windows, Open the `CANDLEHOME\TMAITM6_x64\KUDENV_<instance_name>` file. And edit the line, `KUD_DB2_CLIENT_INST` as `KUD_DB2_CLIENT_INST=<client instance name under which remote Db2 server instance is cataloged>`
5. Run the following command to start the agent:

```
Linux | UNIX | install_dir/bin/db2-agent.sh start instance_name
Windows | install_dir\bin\db2-agent.bat start instance_name
```

**Remember:** While monitoring remote Db2 server instance from UNIX or Linux, the command must be executed with the client instance owner under which remote server instance is cataloged.

### What to do next

- Grant privileges to the Db2 agent user to view data for some attributes of the Db2 agent. For information about granting these privileges, see [“Granting privileges for viewing Db2 agent metrics” on page 287](#).
- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## Granting privileges for viewing Db2 agent metrics

To monitor the Db2 agent resources, a Db2 user must have the Db2 agent SYSADM, SYSCTRL, SYSMAINT, and SYSMON authorities for the monitored instance to view the data for some attributes of Db2.

### About this task

To view the monitoring data that the agent collects for all the attributes on the dashboard, the Db2 user must have specific privileges. To assign these privileges to the Db2 user, run the script file that is present at the following location:

```
Linux | UNIX | install_dir/config/KudGrantUserPermissions.sh
Windows | install_dir\TMAITM6_x64\KudGrantUserPermissions.bat
```

A Db2 user with the SYSADM authority can run the script to grant privileges to itself or to any other Db2 user. For a Db2 instance, use the instance owner, which already has the SYSADM authority, to run the script to grant other permissions to itself or to grant all the permissions to any other Db2 user.

## Procedure

1. For local monitoring, follow these steps.
  - a) On the system where the Db2 is installed, open the Db2 command-line interface.
  - b) Run the following command where *instance\_name* is the name of the Db2 instance and *username* is the name of the Db2 user:

```
Linux | UNIX | install_dir/config/KudGrantUserPermissions.sh
instance_name username
Windows | install_dir\TMAITM6_x64\KudGrantUserPermissions.bat
instance_name username
```

**Note:** For Windows systems, *username* is optional in the command. If a user name is not specified in the command, the privileges are assigned to the default user (system).

2. For remote monitoring, follow these steps.
  - a) Copy `KudGrantUserPermissions.sh` for UNIX or Linux and `KudGrantUserPermissions.bat` for Windows from `install_dir/TMAITM6_x64/` from agent workstation to the Db2 machine.
  - b) Run the following command from Db2 instance owner user where *instance\_name* is the name of the Db2 instance and *username* is the name of the Db2 user:

```
Linux | UNIX | ./KudGrantUserPermissions.sh instance_name username
Windows | KudGrantUserPermissions.bat instance_name username
```

**Remember:** For remote Db2 monitoring on Windows, the *username* must be the user name that is provided during the Db2 agent configuration at client workstation.

## Configuring local environment variables

You can configure local environment variables to change the behavior of the Db2 agent.

### Procedure

1. For Windows, click **Start > All Programs > IBM Monitoring agents > IBM Application Performance Management**.
2. In the **IBM Application Performance Management** window, from the **Actions** menu, click **Advanced > Edit ENV File**.
3. On Linux or AIX systems, go to the command line and edit the `ud.environment` file from the `install_dir/config` directory. Where, `install_dir` is the agent installation directory.

**Note:** The `ud.environment` file is a hidden file.

4. In the environment variable file, enter the values for the environment variables.

For information about the environment variables that you can configure, see [“Local environment variables”](#) on page 288.

### Local environment variables

You can change the behavior of the Db2 agent by configuring the local environment variables.

### Variables for defining the data collection method for the tablespace data set

To set the method for data collection of the tablespace data set, use the following environment variables:

- **KUD\_T1\_BY\_SQL:** Use this variable to set the method of data collection for the tablespace data set by using SQL queries. To enable data collection by using SQL queries, set the value of this variable as Y. To collect data for the tablespace data set by using the snapshot method, set the value of this variable as N. The default value of this variable is N.

**Important:** To collect data by using SQL queries, the Db2 version must be 9.7, or later. Also, the user who starts the Db2 agent must have the SYSADM authority for all databases.

- **KUD\_T1\_DISABLE:** Use this variable to disable the data collection for the tablespace data set. To enable the data collection for the tablespace data set, set the value of this variable as N. To disable the data collection for the tablespace data set, set the value of this variable as Y. The default value of this variable is N.

#### **Variable for excluding the Caching Facility (CF) nodes from data collection**

To exclude CF nodes from the data collection algorithm in pureScale® environment, use the **DB2\_CF\_PARTITION\_NUMS** variable. Set the CF node number to be excluded as the value of the **DB2\_CF\_PARTITION\_NUMS** variable: **DB2\_CF\_PARTITION\_NUMS=<CF node number>**. For example, **DB2\_CF\_PARTITION\_NUMS=1** excludes CF node 1. For more than one CF node, set the **DB2\_CF\_PARTITION\_NUMS** variable value as a list that uses any special symbol from # . : , ; | @ as delimiter. For example, **DB2\_CF\_PARTITION\_NUMS=12 , 13 , 23 , 34**. No default value is set for this variable.

#### **Variable for limiting data collection for the DB2 Table data set**

To set the maximum number of rows that the Db2 agent must return, while collecting data for the DB2® Table data set, use the **KUD\_TABLE\_NUMBER** environment variable. The default value is 10000.

#### **Variable for setting the reload interval of the customized SQL properties file**

To set the reload time interval (in seconds) for the customized SQL properties file, use the **KUD\_CUS\_SQL\_INTERVAL** variable. The default value is 20 seconds.

#### **Variable for limiting the rows in the data collection for Agent Event data set**

To set the number of rows for data collection of the Agent Event data set, use the **KUD\_AGENT\_EVENT\_CACHE** variable. The Agent Event data set provides detailed information about predefined and triggered events and determines problems with the health of the monitored database. The default value is 50.

#### **Variable for limiting the rows in the data collection for DB2 Log Record data set**

To set the number of rows for data collection of the DB2 Log Record data set, use the **KUD\_DBHISTORY\_MAXROW** variable. The DB2 Log Record data set provides historical information about the Db2 archive log. The default value is 500.

#### **Variables for defining the data collection for the DB2 Diagnostic Log data set**

To set the method for data collection of the DB2 Diagnostic Log data set, use the following environment variables:

- **KUD\_DIAGLOG\_BY\_TABLE:** Use this variable to set the method of data collection for the DB2 Diagnostic Log data set. If the value of this variable is set to Y, then data for the DB2 Diagnostic Log data set is collected by using SQL queries. If the value of this variable is set to N, then data for the DB2 Diagnostic Log data set is collected by parsing the `db2diag.log`. The default value of this variable is Y.

**Important:** To collect data by using SQL queries, the Db2 version must be 10, or later.

- **KUD\_DIAGLOG\_TAILCOUNT:** Use this variable to define the number of lines of the `db2diag.log` file that the Db2 agent parses for collecting data for the DB2 Diagnostic Log data set. This variable limits the Db2 agent to process the Db2 agent log file so that only the latest messages and events are monitored. The default value of this variable is 1000.
- **KUD\_DIAGLOG\_CACHE:** Use this variable to limit the number of log records that are displayed on the dashboard for the DB2 Diagnostic Log data set. The default value of this variable is 20.
- **KUD\_DIAGLOG\_INTERVAL:** Use this variable to define the reload time interval (in seconds) for the `db2diag.log` file for data collection for the DB2 Diagnostic Log data set. The default value of this variable is 30 seconds.

- **KUD\_DISABLE\_DIAGLOG:** Use this variable to disable the data collection for the DB2 Diagnostic Log data set. To enable the data collection for the DB2 Diagnostic Log data set, set the value of this variable as N. To disable the data collection for the DB2 Diagnostic Log data set, set the value of this variable as Y. The default value of this variable is N.

#### **Variable for setting the query timeout interval**

If an SQL query takes a very long time to complete, it affects the performance of the Db2 agent. To set the query timeout interval for the Db2 agent, use the **KUD\_QUERY\_TIMEOUT** variable. Use this variable to define the maximum amount of time (in seconds) that the Db2 agent waits to receive a response for a query that is sent to the Db2 server. The value for this variable must be less than 300 seconds. The default value of this variable is 45 seconds.

#### **Variable for defining the data collection for the DB2 Database01 (Superseded) data set**

The agent must not trigger ASN queries to collect data for the DB2 Database01 (Superseded) data set when ASN schemas are not present. To enable the execution of the ASN queries, use the **KUD\_REPLICATION\_ON** variable. If the value of this variable is set to Y, the Db2 agent runs ASN queries even when the ASN schemas are not present. If the value of this variable is set to N, the Db2 agent does not run the ASN queries. The default value of this variable is Y.

#### **Variable for configuring the monitor switches when collecting data by using the snapshot method**

If you want to collect the Db2 agent monitoring data by using the snapshot method, enable the Db2 monitor switch for the data set. To enable the Db2 monitor switch, use the **KUD\_MON\_SWITCH\_OVERRIDE** variable. The list of Db2 monitor switches is as follows:

##### **LOCK**

Lock Information

##### **SORT**

Sorting Information

##### **STATEMENT**

SQL Statement Information

##### **TABLE**

Table Activity Information

##### **TIMESTAMP**

Take Timestamp Information

##### **UOW**

Unit of Work Information

If the value of this variable is set to Y, the Db2 agent retains the configuration setting of the Db2 monitor switches. If the value of this variable is set to N, the Db2 enables all the monitor switches to collect data. The default value of this variable is N.

#### **Variable for tracing the Db2 snapshot buffer data of an data set**

To view the data that is collected for an data set by using the snapshot method, use the **KUD\_SNAPSHOT\_DUMP** variable. If the value of this variable is set to Y, the Db2 agent dumps the snapshot buffer data for attribute groups in the agent log file. If the value of this variable is set to N, the Db2 agent does not dump the snapshot buffer data in the agent log file. The default value of this variable is N.

#### **Variable for tracing the Db2 agent by using the snapshot buffer data of an data set**

To trace the Db2 agent by using the snapshot buffer data that is collected for an data set, use the **KUD\_SNAPSHOT\_READIN** variable. To enable the tracing of Db2 agent, set the value of this variable as Y. To disable the tracing of Db2 agent, set the value of this variable as N.

### Variable for defining the data collection method for the Locking Conflict data set

To set the method of data collection for the Locking Conflict data set, use the **KUD\_LOCKCONFLICT\_BY\_SQL** variable. To collect data for the Locking Conflict data set by using SQL queries, set the value of this variable as Y. To collect data for the Locking Conflict data set by using the snapshot method, set the value of this variable as N. The default value of this variable is Y.

**Important:** To collect data by using SQL queries, the Db2 version must be 9.7 FP1, or later. Also, the user who starts the Db2 agent must have SYSADM authority for all databases.

### Variable to monitor remote Db2 server on Windows

**KUD\_DB2\_CLIENT\_INST:** Set this variable to Db2 client instance name under which remote Db2 server instance is cataloged. You need to set this variable only if you are using remote monitoring where agent is on Windows.

## Prerequisites for Remote Monitoring

You can use Monitoring Agent for Db2 for remote monitoring. Refer the topic for prerequisites of remote monitoring of Db2.

### About this task

For remote monitoring of Db2, you must first do the basic Db2 client/server environment setup. Do this setup for Windows and UNIX or Linux.

For this set up a user must have Db2 SYSADM or SYSCTRL authority.

**Remember:** Run all the steps on agent workstation except for step 2.

### Procedure

1. On the Db2 agent workstation, install Db2 client. The version of this client must be greater than or equal to that of Db2 server instance version that is to be monitored.
2. Verify that the communication protocol for Db2 instance is TCPIP.
  - a) To verify, run the command **db2set** on the Db2 command line.
  - b) If it is not set to TCPIP, then run **db2set DB2COMM=tcPIP** in Db2 command line.

**Important:** This step is done at the server side.

3. Catalog the remote server instance at Db2 agent workstation with following command.

**Important:** The server instance is to be cataloged under the client instance. So run following command on the client instance.

```
db2=>CATALOG TCPIP NODE<node_name> REMOTE <hostname/ip_address> SERVER <service_name/  
port_number>
```

on Db2 where

- a. *<node\_name>* represents a local nickname of Db2 instance on client component.

**Note:** For UNIX or Linux, *<node\_name>* must not be same as of any Db2 client or Db2 server instance name available on the same workstation.

- b. *<hostname/ip\_address>* represents name or IP address of the Db2 server workstation.
- c. *<service\_name/port\_number>* at which Db2 TCPIP configured.

To catalog Db2 server instance running on port number 50000 on remote server "**myserver**" as node "db2node", enter the following command from a Db2 command line

```
db2 => CATALOG TCPIP NODE db2node REMOTE myserver SERVER 50000
```

For more details on catalog node, refer [Cataloging a TCP/IP node from a client using the CLP](#)

4. If Db2 agent workstation is UNIX/Linux,

- Create a user with node name, which is used in cataloging command

Issue the command

```
useradd -g <group> -m -d <home_dir> <user> -p <password>
```

where

- **<group>** represents a group for the DB2 UDB instance owners.
  - **<user>** represents a local **username** on client workstation. **Username** must be same as node name by which the server instance has been cataloged on agent machine.
- Check the Db2 client instance name under which remote Db2 server instance is cataloged and assign the read, write, execute permissions of the newly created user's home directory to the owner of this instance. This step is necessary to make the client Db2 environment available for operations on remote node
  - Issue the command

```
chmod -R 775 /home/<nodename>
```

where

- **<nodename>** represents a local username of Db2 instance on client component

5. Catalog all the databases that you want to monitor on the client instance present at Db2 agent workstation.

Issue the command in the Db2 CLP to catalog the database.

```
CATALOG DATABASE <db_name> AS <db_alias> AT NODE <node_name>authentication server
```

- a. **<db\_name>** represents server database name.
- b. **<db\_alias>** represents local nickname for database at Db2 client.
- c. **<node\_name>** represents a local nickname of Db2 instance on client component at which database is cataloged.

To catalog a database called "sample" on catalog node "db2node" with alias as "dbAlias1", enter the following command from a Db2 prompt.

```
db2 => CATALOG DATABASE sample AS dbAlias1 AT NODE db2node authentication server
```

## Configuring Hadoop monitoring

You must configure the Monitoring Agent for Hadoop so that the agent can collect data of a Hadoop cluster that it monitors. The agent can monitor a single node Hadoop cluster and a multi-node Hadoop cluster.

### Before you begin

Review the hardware and software prerequisites. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\) for the Hadoop agent](#).

The IBM Cloud App Management Hadoop agent is installed with IBM Cloud App Management Extension Packs. Install the Hadoop agent with the extension pack and proceed to configure the agent. For more information about extension packs for Hadoop, see [“Part numbers” on page 71](#).

Ensure that the following hosts can be resolved from the computer where the Hadoop agent is installed:

- All the Hadoop hosts that you want to configure, such as NameNode, ResourceManager, and so on
- Hadoop hosts with only NodeManager role

For example, you can complete these steps to resolve hosts:



- Add the IP address, host name, and fully qualified domain name of all the Hadoop hosts to the hosts file that is available at the following path:

- **Windows** C:\Windows\System32\drivers\etc\hosts
- **Linux** **AIX** /etc/hosts

- Add the computer where the Hadoop agent is installed in the same domain as that of Hadoop hosts.

**Remember:** To monitor a Hadoop cluster that is secured with Kerberos SPNEGO-based authentication, ensure that all the hosts can be resolved from the computer where the Hadoop agent is installed.

### About this task

The Hadoop agent is a single instance agent. You must configure the agent manually after it is installed. The Hadoop agent can be configured on Windows, Linux, and AIX systems.

### Remember:

- For a single node Hadoop cluster, the same node performs all the roles, such as NameNode, ResourceManager, and secondary NameNode according to configuration of the Hadoop cluster. However, for a multi-node Hadoop cluster, different Hadoop nodes perform these roles.
- When you configure the agent, the agent automatically detects DataNodes and NodeManagers in the Hadoop cluster that is being monitored.

Complete the configuration steps that are specified in the subsequent topics. Ensure that you specify the host names according to the following guidelines when you configure the agent.

- The host name of various daemon processes (NameNode, ResourceManger, and so on) that you specify must be the same (case and format) as the host names that are configured for the socket-based agent.
- The fully qualified domain name (FQDN) must be used when you specify a host name. For example, `hos1.ibm.com`. If the length of the FQDN exceeds 25 characters, specify only the short host name without the domain name. For example, if the FQDN of a host is `myhadoopclustersetupnode.ibm.com`, the short host name is `myhadoopclustersetupnode`.

After you configure the agent that is upgraded, and view data in the Cloud App Management console, revert the changes that were made in the `hadoop-metrics2.properties` file for the Hadoop agent. For details, see [“Upgrading your ICAM Agents” on page 1387](#).

On Windows systems, you can run the Hadoop agent with a non-administrator user. However, such user requires a specific permission to view data in the dashboards. For information about how to grant this permission, see [“Granting permission to non-admin users” on page 300](#).

The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history” on page 52](#).

## Configuring the agent on Windows systems

You can configure the agent on Windows systems by using the **IBM Performance Management** window.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Monitoring Agent for Hadoop**.
3. Click **Configure agent**.



**Attention:** If **Configure agent** is disabled, click **Reconfigure**.

The **Configure Monitoring Agent for Hadoop** window opens.

4. To monitor the Hadoop cluster with the Kerberos SPNEGO-based authentication enabled, complete these steps:

- a) Under **Is Kerberos SPNEGO-based authentication for HTTP based Hadoop services in Hadoop cluster enabled**, click **Yes**.  
If you do not have Kerberos SPNEGO-based authentication to secure REST endpoints of HTTP based Hadoop services in the Hadoop cluster, click **No** and then the values for the **Realm name**, **KDC Hostname**, **SPNEGO principal name** and **SPNEGO keytab file** fields can be kept as blank.
  - b) In the **Realm name** field, enter the name of the Kerberos realm that is used to create service principals.  
Usually, a realm name is the same as your domain name. For instance, if your computer is in the `tivoli.ibm.com` domain, the Kerberos realm name is `TIVOLI.IBM.COM`. This name is case sensitive.
  - c) In the **KDC Hostname** field, enter the fully qualified domain name (FQDN) of the Key Distribution Center (KDC) host for the specified realm.  
You can also specify the IP address of the KDC host instead of FQDN. In case of Active Directory KDC, Domain controller is the KDC host.
  - d) In the **SPNEGO principal name** field, enter the name of the Kerberos principal that is used to access SPNEGO authenticated REST endpoints of HTTP-based services.  
The name is case sensitive, and the name format is `HTTP/fully_qualified_host_name@kerberos_realm`
  - e) In the **SPNEGO keytab file** field, enter the name of the keytab file for the SPNEGO service with its full path, or click **Browse** and select it.  
The keytab file contains the names of Kerberos service principals and keys. This file provides direct access to Hadoop services without requiring a password for each service. The file can be located at the following path: `etc/security/keytabs/`  
Ensure that the SPNEGO principal name and the keytab file belong to the same host. For instance, if the principal name is `HTTP/abc.ibm.com@IBM.COM`, the keytab file that is used must belong to the `abc.ibm.com` host.  
If the agent is installed on a remote computer, copy the keytab file of the principal to the remote computer at any path, and then specify this path in the **SPNEGO keytab file** field.
  - f) Click **Next**.
5. To monitor Hadoop cluster with HTTPS/SSL enabled, complete these steps:
    - a) Under **Is Hadoop Cluster SSL enabled**, click **Yes**  
If you do not want the SSL enabled Hadoop cluster select **No** and then the values for the **TrustStore file path**, **TrustStore Password** fields can be kept as blank.
    - b) In **TrustStore file path**, select the TrustStore file stored at your local machine.  
This file can be copied from the Hadoop cluster to your local machine and then used for configuration.
    - c) In **TrustStore Password**, enter the password you created while configuring the TrustStore file.
  6. To specify values for the parameters of the Hadoop cluster, complete these steps:
    - a) In the **Unique Hadoop Cluster Name** field, enter the unique name for the Hadoop cluster indicating Hadoop version and flavor. The maximum character limit for this field is 12.
    - b) In the **NameNode Hostname** field, enter the host name of the node where the daemon process for NameNode runs.
    - c) In the **NameNode Port** field, enter the port number that is associated with the daemon process for NameNode. The default port number is 50070.
    - d) In the **ResourceManager Hostname** field, enter the host name of the node where the daemon process for ResourceManager runs.
    - e) In the **ResourceManager Port** field, enter the port number that is associated with the daemon process for ResourceManager. The default port number is 8088.
    - f) Optional: In the **JobHistoryServer Hostname** field, enter the host name of the node where the daemon process for JobHistoryServer runs.

- g) Optional: In the **JobHistoryServer Port** field, enter the port number that is associated with the daemon process for JobHistoryServer. The default port number is 19888.
  - h) Optional: In the **Additional NameNode Hostname** field, enter the host name where the daemon process for a Standby NameNode or a Secondary NameNode runs.
  - i) Optional: In the **Additional NameNode Port** field, enter the port number that is associated with the daemon process for a Standby NameNode or a Secondary NameNode.
 

**Remember:** If the additional NameNode is a Standby NameNode, the default port number that is associated with the Standby NameNode daemon process is 50070. If the additional NameNode is a Secondary NameNode, the default port number that is associated with the Secondary NameNode daemon process is 50090.
  - j) Click **Test Connection** to verify connection to the specified host names and ports.
 

After you click **Test Connection**, an appropriate validation message is displayed when:

    - The connection to the specified host names and ports is made or failed.
    - A value for a host name is kept as blank.
    - A value for a port is kept as blank.
    - A non-integer value is specified for a port number.

Update the configuration values as suggested in the validation messages, and verify the connection again.
  - k) Optional: To add Standby ResourceManagers in the Hadoop cluster, click **Yes** under **Standby ResourceManager (s) in Hadoop Cluster**.
 

You are prompted to add the details of Standby ResourceManagers later.
  - l) Optional: To monitor Hadoop services in the Hadoop cluster that is managed by Apache Ambari, click **Yes** under **Monitoring of Hadoop services for Ambari based Hadoop installations**, and then click **Next**.
7. Optional: To specify the details of the Ambari server for monitoring Hadoop services, complete the following steps:
- a) In the **Ambari server Hostname** field, enter the host name where the Ambari server runs.
  - b) In the **Ambari server Port** field, enter the port number that is associated with the Ambari server. The default port number is 8080.
  - c) In the **Username of Ambari user** field, enter the name of the Ambari user.
  - d) In the **Password of Ambari user** field, enter the password of the Ambari user.
  - e) Click **Next**.
8. To specify values for the Java parameters, complete these steps:
- a) From the **Java trace level** list, select a value for the trace level that is used by Java providers.
  - b) Optional: In the **JVM arguments** field, specify a list of arguments for the Java virtual machine. The list of arguments must be compatible with the version of Java that is installed along with the agent.
  - c) Click **Next**.
9. Optional: To add Standby ResourceManagers, complete the following steps:
- a) Click **New**.
  - b) In the **Standby ResourceManager Hostname** field, enter the host name of the node where the daemon process for Standby ResourceManager runs.
  - c) In the **Standby ResourceManager Port** field, enter the port number that is associated with the daemon process for Standby ResourceManager. The default port number is 8088.
  - d) Click **Test Connection** to validate connection to the specified host name and the port number.
 

After you click **Test Connection**, an appropriate validation message is displayed when:

    - The connection to the specified host names and ports is made or failed.

- A value for a host name is kept as blank.
- A value for a port is kept as blank.
- A non-integer value is specified for a port number.

Update the configuration values as suggested in the validation messages, and verify the connection again.

- e) Repeat steps [a](#), [b](#), and [c](#) to add more Standby ResourceManagers.  
If you want to remove any of the Standby ResourceManagers, click **Delete** corresponding to the Standby ResourceManager that you want to remove.
  - f) Click **Next**.
10. In the **Class path for external jars** field, specify the class path for JAR files.  
This class path is added to the class path that is generated by the agent. You can keep this field blank.
  11. Click **OK**.  
The specified configuration settings are saved.
  12. Right-click **Monitoring Agent for Hadoop** and click **Start**.

### What to do next



1. Enable the subnode events to view eventing thresholds of the Hadoop agent. For information about enabling subnode events, see [“Configuring the dashboard for viewing Hadoop events”](#) on page 300.
2. Log in to the Cloud App Management console to view data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring the agent on Linux and AIX systems

To configure the agent on Linux and AIX systems run the configuration script and respond to prompts.

### Procedure

1. On the command line, run the following command: **install\_dir/bin/hadoop-agent.sh config**.  
Where *install\_dir* is the installation directory of Hadoop agent.  
The agent is installed at the following default installation directory:  
`/opt/ibm/apm/agent`
2. When the command line displays the following message, type 1 to continue with the configuration steps and press Enter.  
Edit "Monitoring Agent for Hadoop" setting? [1= yes, 2= No]
3. When the command line displays the following message, type 1 to specify values for monitoring the Hadoop cluster with the Kerberos SPNEGO-based authentication enabled, and press **Enter**.  
Otherwise, type 2 and press **Enter**, and you can keep a blank value for the **Realm name**, **KDC Hostname**, **SPNEGO principal name**, and **SPNEGO keytab file** fields:  
Is Kerberos SPNEGO-based authentication for HTTP based Hadoop services in Hadoop cluster enabled\?: [ 1=Yes, 2=No (default is: 2)]
  - a) For the **Realm name** parameter, enter the name of the Kerberos realm that is used to create service principals.  
Usually, a realm name is the same as your domain name. For instance, if your computer is in the `tivoli.ibm.com` domain, the Kerberos realm name is `TIVOLI.IBM.COM`. This name is case sensitive.
  - b) In the **KDC Hostname** field, enter the fully qualified domain name (FQDN) of the Key Distribution Center (KDC) host for the specified realm. You can also specify the IP address of the KDC host instead of FQDN. In case of Active Directory KDC, Domain controller is the KDC host

- c) For the **SPNEGO principal name** parameter, enter the name of the Kerberos principal that is used to access SPNEGO authenticated REST endpoints of HTTP-based services.  
The name is case sensitive, and the name format is `HTTP/fully_qualified_host_name@kerberos_realm`
- d) For the **SPNEGO keytab file** parameter, enter the name of the keytab file for the SPNEGO service with its full path.  
The keytab file contains the names of Kerberos service principals and keys. This file provides direct access to Hadoop services without requiring a password for each service. The file can be located at the following path: `etc/security/keytabs/`  
Ensure that the SPNEGO principal name and the keytab file belong to the same host. For instance, if the principal name is `HTTP/abc.ibm.com@IBM.COM`, the keytab file that is used must belong to the `abc.ibm.com` host.  
If the agent is installed on a remote computer, copy the keytab file of the principal to the remote computer at any path, and then specify this path for the **SPNEGO keytab file** parameter.
4. When the command line displays the following message, type 1 to specify values for monitoring the Hadoop cluster with the SSL enabled, and press **Enter**. Otherwise, type 2 and press **Enter**, and you can keep a blank value for the **TrustStore file path** and **TrustStore Password** fields:  
Is Hadoop Cluster SSL enabled [ 1=Yes, 2=No (default is: 2)
- a) In **TrustStore file path**, specify the path of TrustStore file stored at your local machine.  
This file can be copied from the Hadoop cluster to your local machine and then used for configuration.
- b) In **TrustStore Password**, specify the password you created while configuring the TrustStore file.
5. When you are prompted to enter the details of the Hadoop cluster, specify an appropriate value for each of the following parameters, and press Enter.
- a) In the **Unique Hadoop Cluster Name**, specify the unique name for the Hadoop cluster indicating Hadoop version and flavor. The maximum character limit for this field is 12.
- b) For the **NameNode Hostname** parameter, specify the host name of the node where the daemon process for NameNode runs, and press Enter.
-  **Attention:** If you press Enter without specifying a host name, you are prompted to enter the host name.
- c) For the **NameNode Port** parameter, specify the port number that is associated with the daemon process for NameNode, and press Enter. The default port number is 50070.
- d) For the **ResourceManager Hostname** parameter, specify the host name of the node where the daemon process for ResourceManager runs, and press Enter.
-  **Attention:** If you press Enter without specifying a host name, you are prompted to enter the host name.
- e) For the **ResourceManager Port** parameter, enter the port number that is associated with the daemon process for ResourceManager. The default port number is 8088.
6. Optional: When you are prompted to add the details of the following parameters of the Hadoop cluster, accept the default value or specify an appropriate value for each of the following parameters, and press Enter:
- a) For the **JobHistoryServer Hostname** parameter, enter the host name of the node where the daemon process for JobHistoryServer runs.
- b) For the **JobHistoryServer Port** parameter, enter the port number that is associated with the daemon process for JobHistoryServer. The default port number is 19888.
- c) For the **Additional NameNode Hostname** parameter, enter the host name of the node where the daemon process for a Secondary or a Standby NameNode runs.

- d) For the **Additional NameNode Port** parameter, enter the port number that is associated with the daemon process for a Secondary or a Standby NameNode. The default port number for a Secondary NameNode is 50090. For a Standby NameNode, the default port number is 50070.
7. Optional: When the command line displays the following message, enter 1 to add details of Standby ResourceManagers for high-availability cluster, and press Enter.  
Standby ResourceManager(s) in Hadoop Cluster [ 1=Yes, 2=No ] (default is: 2):
8. When the command line displays the following message, specify 1 and press Enter to monitor Hadoop services in the Hadoop cluster that is managed by Ambari:  
Monitoring of Hadoop services for Ambari based Hadoop installations [ 1=Yes, 2=No ] (default is: 2):  
Otherwise, retain the default value of 2 and press Enter. If you enable the monitoring of Hadoop services, specify a value for each of the following parameters of Ambari server, and press Enter:
- a) For the **Ambari server Hostname** parameter, enter the host name where the Ambari server runs.
- b) For the **Ambari server Port** parameter, enter the port number that is associated with the Ambari server.  
The default port number is 8080.
- c) For the **Username of Ambari user** parameter, enter the name of the Ambari user.
- d) For the **Password of Ambari user** parameter, enter the password of the Ambari user.
9. When the command line displays the following message, select the appropriate Java trace level and press Enter:  
This parameter allows you to specify the trace level used by the Java providers Java trace level [ 1=Off, 2=Error, 3=Warning, 4=Information, 5=Minimum Debug, 6=Medium Debug, 7=Maximum Debug, 8=All ] (default is: 2)
10. Optional: When the command line displays the following message, specify the arguments for the Java virtual machine, and press Enter. The list of arguments must be compatible with the version of Java that is installed along with the agent.  
This parameter allows you to specify an optional list of arguments to the java virtual machine JVM arguments (default is:)
11. Optional: When the command line displays the following message, enter 1 to add the following details of Standby ResourceManagers, and press Enter:  
Edit "Hadoop High Availability(HA) Cluster with Standby ResourceManagers" settings, [1=Add, 2=Edit, 3=Del, 4=Next, 5=Exit] (default is: 5): 1
- a) For the **Standby ResourceManager Hostname** parameter, enter the host name of the node where the daemon process for Standby ResourceManager runs.
- b) For **Standby ResourceManager Port**, enter the port number that is associated with the daemon process for Standby ResourceManager. The default port number is 8088.
- c) When you are prompted, enter 1 to add more Standby ResourceManagers, and repeat steps [a](#) and [b](#), or enter 5 to go to the next step.
- To edit the configuration settings of a specific Standby ResourceManager, type 4 and press Enter until you see the host name of the required Standby ResourceManager.
  - To remove a Standby ResourceManager, type 3 and press Enter after you see the host name of the Standby ResourceManager that you want to remove.
12. When you are prompted, enter the class path for the JAR files that the Java API data provider requires, and press Enter.  
The specified configuration values are saved, and a confirmation message is displayed.
13. Run the following command to start the agent: **install\_dir/bin/hadoop-agent.sh start**

## What to do next

1. Enable the subnode events to view eventing thresholds of the Hadoop agent. For information about enabling subnode events, see [“Configuring the dashboard for viewing Hadoop events” on page 300.](#)

2. Log in to the Cloud App Management console to view data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring the agent by using the silent response file

You can use the silent response file to configure the Hadoop agent on Linux, AIX, and Windows systems.

### About this task

The silent response file contains the agent configuration parameters. For some parameters, the default values are provided in comments. You can specify different values for these parameters, and remove the comment tags that are placed at the beginning of the parameters.

### Procedure

1. Open the silent response file that is available at this path: `install_dir\samples\hadoop_silent_config.txt`
2. In the response file, complete the following steps:
  - a) When you want to monitor the Hadoop Cluster that is enabled for Kerberos SPNEGO-based authentication, type yes and enter values for the following parameters:

```
HADOOP_REALM_NAME
HADOOP_KDC_HOSTNAME
HADOOP_PRINCIPAL_NAME
HADOOP_SPNEGO_KEYTAB
```

- b) When you want to monitor the Hadoop Cluster that is SSL enabled, type yes and enter values for the following parameters:

```
HADOOP_TRUSTSTORE_PATH
HADOOP_TRUSTSTORE_PASSWORD
```

- c) Enter values for the following parameters of Cluster, NameNode (NN), ResourceManager (RM), and Job History Server (JHS):

```
HADOOP_CLUSTER_NAME (optional)
HADOOP_NN_HOSTNAME
HADOOP_NN_PORT
HADOOP_RM_HOSTNAME
HADOOP_RM_PORT
HADOOP_JHS_HOSTNAME (optional)
HADOOP_JHS_PORT (optional)
```

- d) Optional: For the **HADOOP\_ADDITIONAL\_NN\_HOSTNAME** parameter, specify the host name of the Standby or Secondary NameNode.
- e) Optional: For the **HADOOP\_ADDITIONAL\_NN\_PORT** parameter, specify the port number of the Standby or Secondary NameNode.

**Remember:** If the additional NameNode is a Standby NameNode, the default port number that is associated with the Standby NameNode daemon process is 50070. If the additional NameNode is a Secondary NameNode, the default port number that is associated with the Secondary NameNode daemon process is 50090.

- f) Optional: For the **Hadoop\_SRM** parameter, type Yes to add Standby ResourceManagers for a high-availability cluster, and go to [step g](#).
- g) Optional: To monitor Hadoop services in the Hadoop cluster that is managed by Ambari, enter values for each of the following parameters, and press Enter:

```
AMBARI_SERVER_HOSTNAME
AMBARI_SERVER_PORT
USERNAME_OF_AMBARI_USER
PASSWORD_OF_AMBARI_USER
```

- h) For the **JAVA\_TRACE\_LEVEL** parameter, specify the appropriate trace level.

- i) Optional: For the **JAVA\_JVM\_ARGS** parameter, specify arguments for the Java™ virtual machine.
- j) Optional: Add the host name and the port number of a Standby ResourceManager in the following format: `HADOOP_SRM_PORT.hadoop_srm_config_sec_1=8088`  
Where, `hadoop_srm_config_sec_1` is the host name of the node where the daemon process for Standby ResourceManager runs, and 8088 is the default port number. To add more Standby ResourceManagers, add the host name and port number of other Standby ResourceManagers on new lines in the same format.

3. Save the response file, and run the following command:

```
Linux | UNIX install_dir/bin/hadoop-agent.sh config install_dir/
samples/hadoop_silent_config.txt
Windows install_dir/bin/hadoop-agent.bat config install_dir/samples/
hadoop_silent_config.txt
```

4. Start the agent:

```
Linux | UNIX Run the following command: install_dir\bin\hadoop-agent.sh
start
Windows Right-click Monitoring Agent for Hadoop and then click Start.
```

### What to do next

1. Enable the subnode events to view eventing thresholds of the Hadoop agent. For information about enabling subnode events, see [“Configuring the dashboard for viewing Hadoop events”](#) on page 300.
2. Log in to the Cloud App Management console to view data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Granting permission to non-admin users

On Windows systems, grant the *Debug program* permission to a non-admin user for running the Hadoop agent. This permission is required to view data in the Hadoop agent dashboards.

### Procedure

Complete the following steps on the system where the Hadoop agent is installed:

1. Click **Start > Control Panel > Administrative Tools**.
2. Double-click **Local Security Policy**.
3. In the Security Settings pane, expand **Local Policies** and click **User Rights Assignment**.
4. Right-click **Debug programs** and click **Properties**.
5. Click **Add User or Group**, and add the non-admin user name to which you want to grant this permission.
6. Click **OK**.

### What to do next

Configure and run the Hadoop agent with the non-admin user.

## Configuring the dashboard for viewing Hadoop events

You must configure the dashboard to enable the subnode events so that the **Events** tab can display Hadoop events.

### About this task

The default value for **Enable Subnode Events** is false. Change this value to true for viewing Hadoop events.



## Procedure

1. Open the Cloud App Management console and go to **System Configuration**.
2. On the **Advanced Configuration** page, click **UI Integration** under **Configuration Categories**.
3. From the **Enable Subnode Events** list, select **True**.
4. Click **Save**.

## Configuring HTTP Server agent monitoring

The HTTP Server agent starts automatically after installation. To enable data collection, make sure that the HTTP server is running, and edit the HTTP Server configuration file so that it includes a reference to the HTTP Server agent data collector configuration file..

### Before you begin

The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see ["Change history" on page 52](#).

There are two files involved in the configuration of the HTTP Server agent. To view samples of these files, see [Samples](#). Locate and review the following files:

### The HTTP Server agent data collector configuration file

After you install the HTTP Server agent, it discovers the HTTP server and generates a data collector configuration file in the `install_dir/tmp/khu` directory where `install_dir` is the directory where the HTTP Server agent is installed.

If you have multiple HTTP Servers in your environment, one HTTP Server agent configuration file is generated per HTTP server.

The HTTP Server agent configuration file name is composed of two parts and has following format:

```
khu.full path of the HTTP Server configuration file name.conf
```

The first part of the agent configuration file name is `khu`, in which `hu` is the HTTP server agent code. The second part of the agent configuration file name is created by using the full path and name of the HTTP server configuration file, in which `/` is replaced by `.` . For example, possible file names are as follows:

```
Linux | UNIX | khu.usr.local.apache24.conf.httpd.conf
```

```
Windows | khu.C:\Program Files\IBM\HTTPServer\conf\httpd.conf
```

The HTTP Server agent data collector configuration file contains the following elements:

- Details about the path of the `httpd.conf` file that the HTTP Server uses, for example, `KhuShmemPath "/IBM/HTTPServer/conf/httpd.conf"`.
- Location of the library to load
- Permissions that are associated with the shared memory

### The HTTP server configuration file

Each HTTP server has a configuration file that by default is called `http_server_install_dir/conf/httpd.conf`, where `http_server_install_dir` is the directory where the HTTP Server is installed. In some environments, this file name might be customized. Check the exact file name with the HTTP server administrator.

## Procedure

1. To activate data collection, you must reference the data collector configuration file in the HTTP server configuration file by using the Include statement. Append the following statement to the end of the HTTP Server configuration file:

```
Include "install_dir/tmp/khu/khu.full path of the HTTP Server configuration file name.conf"
```

For example,

**Linux** | **UNIX** If you have an IBM HTTP Server that is installed in the /opt/IBM/HTTPServer directory and the data collector configuration file is in the following directory:

```
/opt/ibm/apm/agent/tmp/khu/khu.opt.IBM.HTTPServer.conf.httpd.conf
```

Append the following statement to the /opt/IBM/HTTPServer/conf/httpd.conf HTTP server configuration file:

```
Include "/opt/ibm/apm/agent/tmp/khu/khu.opt.IBM.HTTPServer.conf.httpd.conf"
```

**Windows** If you have an IBM® HTTP Server that is installed in the C:\ProgramFiles\IBM\HTTPServer directory and the data collector configuration file is in the following directory:

```
C:\IBM\APM\tmp\khu\khu.C.Program Files.IBM.HTTPServer.conf.httpd.conf
```

Append the following statement to the C:\Program Files\IBM\HTTPServer\conf\httpd.conf HTTP server configuration file:

```
Include "C:\IBM\APM\tmp\khu\khu.C.Program Files.IBM.HTTPServer.conf.httpd.conf"
```

2. Change to the following directory:

```
HTTP_server_installation_directory/bin
```

3. Restart the HTTP Server. For example:

**Linux** | **UNIX**

```
./apachectl -k stop  
./apachectl -k start
```

**Windows**

```
httpd.exe -k stop  
httpd.exe -k start
```

## Results

You have successfully configured the agent.

## What to do next

Now, you can verify the HTTP Server agent data is displayed in the console.

## HTTP Server agent code samples

There are two files involved in the configuration of the HTTP Server agent. They are the HTTP Server agent data collector configuration file and the HTTP server configuration file. A sample for the Instance alias mapping file is also provided to help explain how alias works.

### HTTP Server agent data collector file samples

For IBM HTTP Server version 8 and later, 64-bit, the HTTP Server agent data collector configuration file contains this information:

```
#
# Settings for Monitoring Agent for HTTP Server module.
#

LoadModule khu_module "/tmp/ihs/lx8266/hu/lib/khuapache22dc_64.so"

<IfModule mod_khu.c>
    KhuShmemPerm 660
    KhuShmemPath "/opt/IBM/IHS/conf/httpd.conf"
    KhuCpsPath "/tmp/ihs/tmp/khu/khu_cps.properties"
</IfModule>

Alias /khu "/tmp/ihs/lx8266/hu/etc"
<Directory "/tmp/ihs/lx8266/hu/etc">
    Order deny,allow
    Allow from all
    #Require all granted
</Directory>

LoadModule wrt_module /tmp/ihs/lx8266/hu/lib/mod_wrt_ap22_64.so
WrtOriginID HU:tivvm09_httpd:HUS
```

For IBM HTTP Server version 7, 32-bit, the configuration file contains this information:

```
#
# Settings for Monitoring Agent for HTTP Server module.
#

LoadModule khu_module "/tmp/ihs/lx8266/hu/lib/khuapache22dc_32.so"

<IfModule mod_khu.c>
    KhuShmemPerm 660
    KhuShmemPath "/opt/IBM/HTTPServer/conf/httpd.conf"
    KhuCpsPath "/tmp/ihs/tmp/khu/khu_cps.properties"
</IfModule>

Alias /khu "/tmp/ihs/lx8266/hu/etc"
<Directory "/tmp/ihs/lx8266/hu/etc">
    Order deny,allow
    Allow from all
    #Require all granted
</Directory>

LoadModule wrt_module /tmp/ihs/lx8266/hu/lib/mod_wrt_ap22.so
WrtOriginID HU:linux_httpd:HUS
```

For Apache version 2.4, 64-bit, the HTTP Server agent configuration file contains this information:

```
#
# Settings for Monitoring Agent for HTTP Server module.
#

LoadModule khu_module "/tmp/ihs/lx8266/hu/lib/khuapache24dc_64.so"

<IfModule mod_khu.c>
    KhuShmemPerm 660
    KhuShmemPath "/usr/local/apache24/conf/httpd.conf"
</IfModule>

Alias /khu "/tmp/ihs/lx8266/hu/etc"
<Directory "/tmp/ihs/lx8266/hu/etc">
    Order deny,allow
```

```
Allow from all
Require all granted
</Directory>

LoadModule wrt_module /tmp/ihs/lx8266/hu/lib/mod_wrt_ap24_64.so
WrtOriginID HU:linux-tzsi_httpd:HUS
```

### Instance alias mapping file sample

```
# Monitoring Agent for HTTP Server instance alias mapping
# INSTANCE: auto discovered by agent. Please do NOT modify.
# ALIAS: alias name for the instance. The name will be displayed in APM UI dashboard. It
# must be unique
# among all instances and it must be less than 10 characters and consist of only
# alphanumeric characters.
#
INSTANCE.1=/usr/local/apache24/conf/httpd.conf
ALIAS.1=httpd

INSTANCE.1=/usr/local/apache24/conf/admin.conf
ALIAS.1=admin
```

## Configuring IBM Integration Bus monitoring

The IBM Integration Bus agent is a multiple instance agent. You must create a first agent instance and start it manually.

### Before you begin

- The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history” on page 52](#).
- Make sure that the system requirements for the IBM Integration Bus agent are met in your environment. For the up-to-date system requirement information, see the [Detailed system requirements report](#) for the IBM Integration Bus agent.

### Procedure

1. Make sure the user ID that will be used to start and stop the IBM Integration Bus agent belongs to the **mqm** and **mqbrkr** user groups.
2. **Windows**  
If IBM MQ (WebSphere MQ) is installed on the Windows system, add the IBM MQ (WebSphere MQ) library path to the **PATH** environment variable. So that the IBM Integration Bus agent can load the required IBM MQ (WebSphere MQ) libraries to start.
  - a) Add the IBM MQ (WebSphere MQ) library path to the beginning of the **PATH** environment variable.  
For example, if the installation path of IBM MQ (WebSphere MQ) is C:\IBM\WMQ75, add C:\IBM\WMQ75\bin to the beginning of the **PATH** environment variable of your Windows system.
  - b) Restart the Windows system for the changes to take effect.
3. Configure the IBM Integration Bus agent by specifying the following configuration parameters. There are also some optional configuration parameters that you can specify for the agent. For detailed instructions, see [“Configuring the IBM Integration Bus agent” on page 305](#).
  - Agent ID
  - The installation directory of integration nodes (brokers) that are to be monitored
  - The 64-bit library path of IBM MQ (WebSphere MQ)
4. Configure IBM Integration Bus to enable the data that you want to monitor. See [“Configuring IBM Integration Bus for data enablement” on page 308](#).

5. If you have enabled snapshot data collection for your integration node (broker), configure the IBM Integration Bus agent not to store any snapshot data. For instructions, see [“Disabling snapshot data collection for the agent”](#) on page 312.

## Configuring the IBM Integration Bus agent

You must assign an instance name to the IBM Integration Bus agent and configure the agent before it can start monitoring your IBM Integration Bus environment.

### Before you begin

- Make sure that the user ID that is used to start and stop the agent belongs to the **mqm** and **mqbrkrs** user groups.
- **Windows** If IBM MQ (WebSphere MQ) is installed on the Windows system, add the IBM MQ (WebSphere MQ) library path to the **PATH** environment variable. So that the IBM Integration Bus agent can load the required IBM MQ (WebSphere MQ) libraries to start.
  1. Add the IBM MQ (WebSphere MQ) library path to the beginning of the **PATH** environment variable.  
For example, if the installation path of IBM MQ (WebSphere MQ) is C:\IBM\WMQ75, add C:\IBM\WMQ75\bin to the beginning of the **PATH** environment variable of your Windows system.
  2. Restart the Windows system for the changes to take effect.
- You might need to provide the following information according to your environment during the agent configuration. If you do not know the appropriate configuration value to specify, gather the information from the administrator of IBM MQ (WebSphere MQ) and IBM Integration Bus.
  - If IBM MQ (WebSphere MQ) is installed on the same system with the IBM Integration Bus agent, you must provide the 64-bit library path of IBM MQ (WebSphere MQ).
  - If IBM Integration Bus agent will be configured to monitor the integration nodes of IBM Integration Bus V10 or IBM App Connect Enterprise V11, you must provide the installation directory of IBM Integration Bus V10 or IBM App Connect Enterprise V11.
  - If you want the IBM Integration Bus agent to monitor some specific integration nodes (brokers) instead of all on the same system, you must provide the name and installation path of each integration node (broker).

### About this task

The IBM Integration Bus agent is a multiple instance agent; you must create the first instance and start the agent manually.

You can choose to configure the agent with or without interactions on UNIX or Linux systems. On Windows systems, you can configure the agent without interactions only.

- To configure the agent with interaction, run the configuration script and respond to prompts. See [“Interactive configuration”](#) on page 305.
- To configure the agent without interaction, edit the silent response file and then run the configuration script. See [“Silent configuration”](#) on page 306.

### Interactive configuration

#### Procedure

To configure the agent by running the script and responding to prompts, complete the following steps:

1. Enter the following command:

```
install_dir/bin/iib-agent.sh config instance_name
```

where *instance\_name* is the name that you want to give to the agent instance.

2. After you confirm that you want to configure IBM Integration Bus agent, specify the configuration values for general agent settings.
  - a) When prompted for the **Agent Id** parameter, specify a unique alphanumeric string with a maximum length of 8 characters.
 

**Remember:** The specified string must be unique across your environment.
  - b) When prompted for the **IIB version 10 or ACE version 11 Install Directory** parameter, if you want to monitor integration nodes of IBM Integration Bus V10 or IBM App Connect Enterprise V11, specify the installation directory of IBM Integration Bus V10 or IBM App Connect Enterprise V11. For example, `/opt/ibm/mqsi/ace-11.0.0.3`. If you do not want to monitor IBM Integration Bus V10 and IBM App Connect Enterprise V11, press Enter to accept the default.
 

**Remember:** You can specify only one installation directory for the **IIB version 10 or ACE version 11 Install Directory** parameter. If you installed IBM Integration Bus V10 or IBM App Connect Enterprise V11 in different directories and you want to monitor them all, create multiple agent instances and specify one installation directory of IBM Integration Bus V10 or IBM App Connect Enterprise V11 for each agent instance.
3. Optional: Use the **Monitored Broker Settings** section to specify whether you want to use this agent to monitor only some specific integration nodes (brokers).

By default, all integration nodes (brokers) that are running on the same host system as the IBM Integration Bus agent are monitored, as determined by self-discovery. If you want the agent to monitor some specific integration nodes (brokers), specify the name of the integration node (broker) that you want to monitor and set the **Collect Node Data** setting to No, which is the default value, in the **Monitored Broker Settings** section. There can be multiple **Monitored Broker Settings** sections. Each section controls the monitoring settings for one integration node (broker).

**Tip:** You can specify more than one **Monitored Broker Settings** section. When you edit the **Monitored Broker Settings** section, the following options are available:

- Add: Create a **Monitored Broker Settings** section to configure for another integration node (broker).
- Edit: Modify the settings of current **Monitored Broker Settings** section.
- Del: Delete the current **Monitored Broker Settings** section.
- Next: Move to the next **Monitored Broker Settings** section.
- Exit: Exit the **Monitored Broker Settings** configuration.

4. If you confirm that IBM MQ (WebSphere MQ) is installed on the same system, you are prompted for the **WebSphere MQ 64-bit library path** parameter. Press Enter to accept the default value, which is the 64-bit library path of IBM MQ (WebSphere MQ) automatically discovered by the agent. If no default value is displayed, you must provide the 64-bit library path of IBM MQ (WebSphere MQ) before you proceed to the next step. For example, `/opt/mqm8/lib64`.

**Remember:** If your integration nodes (brokers) use different versions of queue managers, specify the latest version of the IBM MQ (WebSphere MQ) 64-bit library path for this parameter.

5. After the configuration completes, enter the following command to start the agent:

```
install_dir/bin/iib-agent.sh start instance_name
```

## Silent configuration

### Procedure

To configure the agent by editing the silent response file and running the script with no interaction, complete the following steps:

1. Open the following agent silent response file in a text editor.
  - **Linux** | **UNIX** `install_dir/samples/iib_silent_config.txt`
  - **Windows** `install_dir\tmaitm6_x64\samples\iib_silent_config.txt`

where *install\_dir* is the agent installation directory. The default installation directory is as follows:

- **Linux** | **UNIX** /opt/ibm/apm/agent
- **Windows** C:\IBM\APM

2. For the **agentId** parameter, specify a unique alphanumeric string with a maximum length of 8 characters as a short identifier for the agent.

**Remember:** The specified string must be unique across your environment.

3. If you want to monitor the integration nodes of IBM Integration Bus version 10 or ACE version 11, specify the installation directory of IBM Integration Bus V10 or IBM App Connect Enterprise V11 for the **defaultWMBInstallDirectory** parameter. For example, C:\Program Files\IBM\IIB\10.0.0.6\ for a Windows system, or /opt/ibm/mqsi/iib-10.0.0.6 for a Linux system.

If you do not want to monitor IBM Integration Bus V10, this parameter is not required because the IBM Integration Bus agent can automatically discover the integration nodes (brokers) of earlier versions.

**Remember:** You can specify only one installation directory for the **defaultWMBInstallDirectory** parameter. If you installed IBM Integration Bus V10 in different directories and you want to monitor them all, create multiple agent instances and specify one installation directory of IBM Integration Bus V10 for each agent instance.

4. Optional: Specify whether you want to use this agent to monitor only some specific integration nodes (brokers).

By default, all integration nodes (brokers) that are running on the same host system as the IBM Integration Bus agent are monitored, as determined by self-discovery. To monitor specific integration nodes (brokers), set the **collectNodeData** and **WMBInstallDirectory** parameters for each integration node (broker) that you want to monitor.

#### **collectNodeData**

Specifies whether node definition data is collected for the monitored integration node (broker).

The syntax is `collectNodeData.brkr_name=NO|YES`, where *brkr\_name* is the name of the integration node (broker).

The default value is NO. It is recommended to use the default value because node definition data is not supported on the Cloud App Management user interface.

#### **WMBInstallDirectory**

The installation directory of the integration node (broker) to be monitored. The syntax is `WMBInstallDirectory.brkr_name=broker_install_dir`, where *broker\_install\_dir* is the installation directory of the integration node (broker) to be monitored.

**Remember:** For a version 10 integration node, the **WMBInstallDirectory** parameter can override the **defaultWMBInstallDirectory** parameter that you set in the previous step.

For example, to monitor only two integration nodes (brokers) that are named BK1 and BK2, set the parameters as follows:

```
collectNodeData.BK1=NO
collectNodeData.BK2=NO
WMBInstallDirectory.BK1=BK1_install_dir
WMBInstallDirectory.BK2=BK2_install_dir
```

5. To monitor brokers that are earlier than IBM Integration Bus V10, specify the 64-bit library path of IBM MQ (WebSphere MQ) for the **WMQLIBPATH** parameter. For example, C:\Program Files\IBM\WebSphere MQ\bin64 for a Windows system, or /opt/mqm8/lib64 for a Linux system.

**Remember:** If your integration nodes (brokers) use different versions of queue managers, specify the latest version of the IBM MQ (WebSphere MQ) 64-bit library path for this parameter.

6. Save and close the agent silent response file, and then enter the following command:

- **Linux** | **UNIX** `install_dir/bin/iib-agent.sh config instance_name path_to_responsefile`

- **Windows** `install_dir\BIN\iib-agent.bat config "instance_name path_to_responsefile"`

where *instance\_name* is the name of the instance that you configure, and *path\_to\_responsefile* is the full path of the silent response file.



**Warning:** On Windows systems, do not include double quotation marks (") that enclose the full path to the silent response file, as this will cause a configuration error.

7. After the configuration completes, enter the following command to start the agent:

- **Linux** | **UNIX**

```
install_dir/bin/iib-agent.sh start instance_name
```

- **Windows**

```
install_dir\bin\iib-agent.bat start instance_name
```

## Results

Now, you can log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

**Remember:** Whenever you update or migrate a monitored integration node (broker), you must restart the IBM Integration Bus agent after the integration node (broker) upgrade or migration.

## Configuring IBM Integration Bus for data enablement

For some data to be available in the Cloud App Management user interface, you must configure IBM Integration Bus to enable the required data collection.

### Before you begin

Make sure that the IBM Integration Bus agent is configured.

**Remember:** Transaction tracking enablement requires you to restart the integration node (broker).

### About this task

Archive statistics and resource statistics can be monitored by the IBM Integration Bus agent only after the data collection is enabled for the integration node (broker).

Decide what type of data that you want to monitor with the IBM Integration Bus agent and complete the following steps according to your needs.

### Procedure

- To enable archive statistics data collection for the integration node (broker), see [“Enabling archive accounting and statistics data collection” on page 308](#).
- To enable resource statistics data for an integration node (broker), see [“Enabling JVM resource statistics” on page 311](#).

## Enabling archive accounting and statistics data collection

### About this task

To enable archive accounting and statistics collection for message flows that belong to the integration node (broker), issue the **mqsichangeflowstats** command from the bin directory of the integration node (broker) installation directory.

**Remember:** Issue the **mqsichangeflowstats** command to the integration node (broker) according to your requirements for monitoring data. It is recommended that you enable only the statistics that you



require because there can be a lot of data and processing when you have many message flows. For more detailed information about the `mqsichangeflowstats` command, refer to IBM Integration Bus documentation.

**Important:** Cloud App Management does not support snapshot accounting and statistics data due to the amount of data and processing required for the set 20 second snapshot interval. Archive data provides the same exact attributes as snapshot data, and is more suitable for the regular production monitoring provided by Cloud App Management. If you have enabled snapshot data collection for the integration node (broker), remember to configure the IBM Integration Bus agent not to store the snapshot data. For instructions, see [“Disabling snapshot data collection for the agent” on page 312.](#)

## Procedure

- To get most data for message flows, issue the following command. This command is recommended because it does not enable the most detailed terminal statistics that provide invocation counts per terminal per node. The terminal level consumes a lot of storage.

```
mqsichangeflowstats BrokerName -a -g -j -c active -t none -n basic -o xml
```

- In ACE version 11, to get most data for message flows, modify the `node.conf.yaml`/`server.conf.yaml` file as follows. These properties are recommended because they do not enable the most detailed terminal statistics that provide invocation counts per terminal per node. The terminal level consumes a lot of storage.

```
Statistics:
# Application message flows will by default inherit Snapshot and Archive values
# set here
Snapshot:
#publicationOn: 'inactive' # choose 1 of : active|inactive, default inactive
# Ensure Events.OperationalEvents.MQ|MQTT
# is set for outputFormat json,xml
#accountingOrigin: 'none' # choose 1 of : none|basic
#nodeDataLevel: 'none' # choose 1 of : none|basic|advanced
#outputFormat: 'usertrace' # comma separated list of :
#csv,bluemix,json,xml,usertrace
#threadDataLevel: 'none' # choose 1 of : none|basic
Archive:
#archivalOn: 'active' # choose 1 of : active|inactive,
# default inactive
# Ensure Events.OperationalEvents.MQ|MQTT
# is set for outputFormat xml
#accountingOrigin: 'none' # choose 1 of : none|basic
#majorInterval: 60 # Sets the interval in minutes at which
#archive statistics are published
#nodeDataLevel: 'basic' # choose 1 of : none|basic|advanced
#outputFormat: 'xml' # comma separated list of : csv,xml,usertrace
#threadDataLevel: 'none' # choose 1 of : none|basic
```

**Note:** If you want to disable this setting, comment out the lines of `archivalOn: 'active'`, `nodeDataLevel: 'basic'`, and `outputFormat: 'xml'`.

- To get all the data supported by the IBM Integration Bus agent, issue the following command:

```
mqsichangeflowstats BrokerName -a -g -j -c active -t none -n advanced -o xml
```

- In ACE version 11, to get all the data supported by the IBM Integration Bus agent, modify the `node.conf.yaml`/`server.conf.yaml` file as follows:

```
Statistics:
# Application message flows will by default inherit Snapshot and Archive values
# set here
Snapshot:
#publicationOn: 'inactive' # choose 1 of : active|inactive, default inactive
# Ensure Events.OperationalEvents.MQ|MQTT
# is set for outputFormat json,xml
#accountingOrigin: 'none' # choose 1 of : none|basic
#nodeDataLevel: 'none' # choose 1 of : none|basic|advanced
#outputFormat: 'usertrace' # comma separated list of :
#csv,bluemix,json,xml,usertrace
#threadDataLevel: 'none' # choose 1 of : none|basic
Archive:
```

```

archivalOn: 'active'      # choose 1 of : active|inactive, default inactive
                          # Ensure Events.OperationalEvents.MQ|MQTT
                          # is set for outputFormat xml
#accountingOrigin: 'none' # choose 1 of : none|basic
#majorInterval: 60       # Sets the interval in minutes at which
                          # archive statistics are published
nodeDataLevel: 'advanced' # choose 1 of : none|basic|advanced
outputFormat: 'xml'      # comma separated list of : csv,xml,usertrace
#threadDataLevel: 'none' # choose 1 of : none|basic

```

**Note:** If you want to disable this setting, comment out the lines of **archivalOn: 'active'**, **nodeDataLevel: 'advanced'**, and **outputFormat: 'xml'**.

- To reduce the amount of data but still reasonably monitor all message flows without further details, issue the following command:

```
mqsichangeflowstats BrokerName -a -g -j -c active -t none -n none -o xml
```

- In ACE version 11, to reduce the amount of data but still reasonably monitor all message flows without further details, modify the `node.conf.yaml/server.conf.yaml` file as follows:

```

Statistics:
# Application message flows will by default inherit Snapshot and Archive values
#set here
Snapshot:
  #publicationOn: 'inactive' # choose 1 of : active|inactive, default inactive
                          # Ensure Events.OperationalEvents.MQ|MQTT
                          # is set for outputFormat json,xml
  #accountingOrigin: 'none' # choose 1 of : none|basic
  #nodeDataLevel: 'none'   # choose 1 of : none|basic|advanced
  #outputFormat: 'usertrace' # comma separated list of :
                          # csv,bluemix,json,xml,usertrace
  #threadDataLevel: 'none' # choose 1 of : none|basic
Archive:
  archivalOn: 'active'    # choose 1 of : active|inactive, default inactive
                          # Ensure Events.OperationalEvents.MQ|MQTT
                          # is set for outputFormat xml
  #accountingOrigin: 'none' # choose 1 of : none|basic
  #majorInterval: 60       # Sets the interval in minutes at which
                          # archive statistics are published
  nodeDataLevel: 'none'   # choose 1 of : none|basic|advanced
  outputFormat: 'xml'    # comma separated list of : csv,xml,usertrace
  #threadDataLevel: 'none' # choose 1 of : none|basic

```

**Note:** If you want to disable this setting, comment out the lines of **archivalOn: 'active'**, **nodeDataLevel: 'none'**, and **outputFormat: 'xml'**.

- If you have a large number of message flows and want to reduce the amount of data, you can specify which message flows to monitor by replacing the `-g` or `-j` option in the previously mentioned commands.
  - To specify a particular integration server (execution group) for enablement, replace `-g` with `-e IntegrationServerName`.
  - To identify a particular message flow for enablement, replace `-j` with `-f MessageFlowName`.
  - If you have grouped your message flows into applications, to specify a particular application for enablement, add `-k ApplicationName` to the `-j` option.
- The IBM Integration Bus agent collects archive accounting and statistics data at the interval of 5 minutes. To set the interval at which the integration node (broker) produces the archive accounting and statistics data to the same interval, issue the following command with the integration node (broker) stopped, and then restart the integration node (broker):

```
mqsichangebroker BrokerName -v 5
```

- In ACE version 11, the IBM Integration Bus agent collects archive accounting and statistics data at the interval of 5 minutes. To set the interval at which the integration node (broker) produces the archive accounting and statistics data to the same interval, modify the `node.conf.yaml/server.conf.yaml` file as follows:

```

Statistics:
# Application message flows will by default inherit Snapshot and Archive values
# set here
Snapshot:
#publicationOn: 'inactive' # choose 1 of : active|inactive, default inactive
# Ensure Events.OperationalEvents.MQ|MQTT
# is set for outputFormat json,xml
#accountingOrigin: 'none' # choose 1 of : none|basic
#nodeDataLevel: 'none' # choose 1 of : none|basic|advanced
#outputFormat: 'usertrace' # comma separated list of :
# csv,bluemix,json,xml,usertrace
#threadDataLevel: 'none' # choose 1 of : none|basic
Archive:
archivalOn: 'active' # choose 1 of : active|inactive, default inactive
# Ensure Events.OperationalEvents.MQ|MQTT
# is set for outputFormat xml
#accountingOrigin: 'none' # choose 1 of : none|basic
majorInterval: 5 # Sets the interval in minutes at which
# archive statistics are published
nodeDataLevel: 'none' # choose 1 of : none|basic|advanced
outputFormat: 'xml' # comma separated list of : csv,xml,usertrace
#threadDataLevel: 'none' # choose 1 of : none|basic

```

## Enabling JVM resource statistics

### About this task

To enable JVM resource statistics for integration servers that belong to the integration node (broker), issue the **mqsichangeresourcestats** command from the bin directory of the integration node (broker) installation directory.

**Remember:** The JVM resource statistics are considered optional because only a few attributes of data are displayed for the high cost of the agent processing this data every 20 seconds. Be sure to consider whether you really need the JVM resource statistics data.

### Procedure

- To enable the statistics across all integration servers in the integration node (broker), issue the following command:

```
mqsichangeresourcestats BrokerName -c active
```

- In ACE version 11, to enable the statistics across all integration servers in the integration node (broker), modify the node.conf.yaml file as follows:

```

Statistics:
# Application message flows will by default inherit Snapshot and Archive values
# set here
Snapshot:
#publicationOn: 'inactive' # choose 1 of : active|inactive, default inactive
# Ensure Events.OperationalEvents.MQ|MQTT
# is set for outputFormat json,xml
#accountingOrigin: 'none' # choose 1 of : none|basic
#nodeDataLevel: 'none' # choose 1 of : none|basic|advanced
#outputFormat: 'usertrace' # comma separated list of :
# csv,bluemix,json,xml,usertrace
#threadDataLevel: 'none' # choose 1 of : none|basic
Archive:
archivalOn: 'active' # choose 1 of : active|inactive, default inactive
# Ensure Events.OperationalEvents.MQ|MQTT
# is set for outputFormat xml
#accountingOrigin: 'none' # choose 1 of : none|basic
majorInterval: 5 # Sets the interval in minutes at which
# archive statistics are published
nodeDataLevel: 'advanced' # choose 1 of : none|basic|advanced
outputFormat: 'xml' # comma separated list of : csv,xml,usertrace
threadDataLevel: 'basic' # choose 1 of : none|basic
Resource:
reportingOn: true # choose 1 of : true|false, default false
.....

```

**Note:** If you want to disable this setting, comment out **reportingOn: true**.

- To enable the statistics for a given integration server in the integration node (broker), issue the following command:

```
mqsichangeresourcestats BrokerName -e IntegrationServerName -c active
```

- In ACE version 11, to enable the statistics for a given integration server in the integration node (broker), modify the `server.conf.yaml` file as follows:

```
Statistics:
# Application message flows will by default inherit Snapshot and Archive values
# set here
Snapshot:
  #publicationOn: 'inactive' # choose 1 of : active|inactive, default inactive
                                # Ensure Events.OperationalEvents.MQ|MQTT
                                # is set for outputFormat json,xml
  #accountingOrigin: 'none' # choose 1 of : none|basic
  #nodeDataLevel: 'none' # choose 1 of : none|basic|advanced
  #outputFormat: 'usertrace' # comma separated list of :
                                # csv,bluemix,json,xml,usertrace
  #threadDataLevel: 'none' # choose 1 of : none|basic
Archive:
  archivalOn: 'active' # choose 1 of : active|inactive, default inactive
                                # Ensure Events.OperationalEvents.MQ|MQTT
                                # is set for outputFormat xml
  #accountingOrigin: 'none' # choose 1 of : none|basic
  majorInterval: 5 # Sets the interval in minutes at which
                    # archive statistics are published
  nodeDataLevel: 'advanced' # choose 1 of : none|basic|advanced
  outputFormat: 'xml' # comma separated list of : csv,xml,usertrace
  threadDataLevel: 'basic' # choose 1 of : none|basic
Resource:
  reportingOn: true # choose 1 of : true|false, default false
```

**Note:** If you want to disable this setting, comment out **reportingOn: true**.

## Disabling snapshot data collection for the agent

Cloud App Management does not support snapshot accounting and statistics data due to the amount of data and processing required for the set 20 second snapshot interval. If you have enabled snapshot data collection for the broker, remember to configure the IBM Integration Bus agent not to store the snapshot data.

### Procedure

- Open the agent configuration file in a text editor. The agent configuration file is in one of the following directories depending on the operating system:

- Linux** `install_dir/config/hostname_qi_instance_name.cfg`
- Windows** `install_dir\TMAITM6_x64\hostname_qi_instance_name.cfg`

where `install_dir` is the agent installation directory; `hostname` is the host name of the operating system; `instance_name` is the agent instance name.

- Edit the file by adding the following parameter to the KqiAgent section:

```
defaultRetainRecentSnapshotSamples=0
```

Example:

```
INSTANCE=inst1 [
SECTION=KqiAgent [ { agentId=inst1 } { instName=inst1 }
{defaultRetainRecentSnapshotSamples=0}
SECTION=MonitorBroker:BRK1 [ { collectNodeData=NO } ]
SECTION=MonitorBroker:BRK2 [ { collectNodeData=NO } ]
]
```

- Save and close the file.
- Restart the IBM Integration Bus agent for the changes to take effect.

## Configuring InfoSphere DataStage monitoring

You must configure the DataStage agent so that the agent can collect data to monitor the health and performance of the DataStage server resources.

### Before you begin

Review the hardware and software prerequisites. For the up-to-date system requirement information, see the Software Product Compatibility Reports (SPCR) for the [Software Product Compatibility Reports for DataStage agent](#).

### About this task

The DataStage agent is a multiple instance agent. You must create the first instance and start the agent manually.

The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history”](#) on page 52.

## Configuring the agent on Windows systems

You can configure the agent on Windows operating systems by using the **IBM Performance Management** window. After you update the configuration values, start the agent to apply the updated values.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Monitoring Agent for InfoSphere DataStage** template, and then click **Configure agent**.

**Remember:** After you configure an agent instance for the first time, the **Configure agent** option is disabled. To configure the agent instance again, right-click on it and then click **Reconfigure...**

3. In the **Enter a unique instance name** field, type an agent instance name and click **OK**.
4. In the **Monitoring Agent for DataStage** window, specify values for the configuration parameters and click **OK**.

For information about the configuration parameters, see [“Configuration parameters of the agent”](#) on page 315.

5. In the **IBM Performance Management** window, right-click the agent instance that you created and click **Start** to start the agent.

## Configuring the agent on Linux systems

To configure the agent on Linux operating systems, you must run the script and respond to prompts.

### Procedure

1. On the command line, change the path to the agent installation directory.  
Example: `/opt/ibm/apm/agent/bin`
2. Run the following command where *instance\_name* is the name that you want to give to the instance:  
`./datastage-agent.sh config instance_name`
3. When the command line displays the following message, type 1 and press enter:  
Edit 'Monitoring Agent for DataStage' setting? [1=Yes, 2=No]
4. Specify values for the configuration parameters when you are prompted.

For information about the configuration parameters, see [“Configuration parameters of the agent”](#) on page 315.

5. Run the following command to start the agent:

```
./datastage-agent.sh start instance_name
```

## Configuring environment variables

You can configure environment variables to change the behavior of the DataStage agent.

### Procedure

1. Open the following file in a text editor:

- a) **Windows** `install_dir\TMAITM6_x64\KDTENV_instance_name`
- b) **Linux** `install_dir/config/.dt.environment`

2. Edit the following environment variables:

- **KDT\_FIRST\_COLLECTION\_INTERVAL**: The time interval in seconds for first data collection. Set this time interval to a duration by which the agent would collect previous Job runs data in the specified time until the agent starts. The default value is 300 seconds (5 minutes). So, if the agent starts at 2:00 PM, it collects the Job runs data from 1:55 PM to 2:00 PM. This is to avoid data storm of historical job runs when the agent starts collecting data. All subsequent agent data collection for job runs fetch only the newly added job runs that took place since the last collection.
- **KDT\_SSL\_CONTEXT**: The SSL protocol that is enabled on the Service Tier (WebSphere Application Server). The default value is TLS.
- **KDT\_META\_SCHEMA\_NAME**: The name of database schema that is created for the metadata repository. The default value is DS0DB for Db2 and xmeta for MSSQL and Oracle databases.
- **KDT\_DATABASE\_SERVICE\_NAME**: The database or service name that is used by the agent to connect to the metadata repository. The default value is XMETA for Db2, xmeta for MSSQL, and ORCL for Oracle databases.
- **KDT\_DISABLED\_ATTRIBUTEGROUP**: A comma-separated list of attribute groups whose data collection needs to be unavailable. Following values can be set as single or multiple for respective attribute group: JobRuns, JobProperties, JobRunLog, JobStages, JobParameters, EngineSystemConfiguration, EngineSystemResources, EngineServiceStatus, EngineStatusSummary, JobActivity, AgentConfiguration, and JobConfiguration.

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

You can use the silent response file to configure the DataStage agent on Linux and Windows system. After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

1. In a text editor, open the silent config file that is available at the following location and specify values for all the parameters:

- Windows** `install_dir\samples\datastage_silent_config.txt`
- Linux** `install_dir\samples\datastage_silent_config_UNIX.txt`
- Windows** `C:\IBM\APM\samples`
- Linux** `/opt/ibm/apm/agent/samples`

For information about the configuration parameters, see [“Configuration parameters of the agent”](#) on page 315.

2. On the command line, change the path to `install_dir\bin`.

3. Run the following command:

```
Windows datastage-agent.bat config instance_name install_dir\samples  
\datastage_silent_config.txt
```

```
Linux datastage-agent.sh config instance_name install_dir\samples  
\datastage_silent_config_UNIX.txt
```

4. Start the agent.

```
Windows In the IBM Performance Management window, right-click the agent instance that you  
created, and click Start.
```

```
Linux Run the following command: ./datastage-agent.sh start instance_name
```

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuration parameters of the agent

While configuring the DataStage agent, you can change the service tier, metadata repository, and advanced configuration parameters.

### Service tier configuration parameters

The configuration parameters that are required for the agent to connect to the service tier.

The following table contains detailed descriptions of the service tier configuration parameters of the DataStage agent.

Parameter name	Description	Mandatory field
Hostname	Hostname of the computer where service tier is installed. If the computer is part of a domain, then provide the fully qualified domain name (FQDN). The default value is localhost.	Yes
HTTPS Port	HTTPS port for REST interface on the computer where service tier is installed. The default value is 9443.	Yes
WAS Username	The user name for connecting to the WebSphere Application Server. The default value is wasadmin.	Yes
WAS Password	The password for connecting to the WebSphere Application Server.	Yes
Confirm WAS Password	The password that is specified in the <b>WAS Password</b> field.	Yes

### Metadata repository configuration parameters

The configuration parameters that are required for the agent to connect to the metadata repository.

The following table contains detailed descriptions of the metadata repository configuration parameters of the DataStage agent.

Parameter name	Description	Mandatory field
Database Type	Database type of the metadata repository. The Db2 default value is 1.	Yes

<i>Table 35. Names and descriptions of the metadata repository configuration parameters (continued)</i>		
<b>Parameter name</b>	<b>Description</b>	<b>Mandatory field</b>
Hostname	Hostname of the computer where metadata repository is installed. If the computer is part of a domain, then provide fully qualified domain name (FQDN). The default value is localhost.	Yes
Database Port	Database port on metadata repository for JDBC Connection. The default value is 50000.	Yes
Database Username	The username for connecting to operations database. The default value is dsodb.	Yes
Database Password	The password for connecting to operations database.	Yes
Confirm Database Password	The password that is specified in the <b>Database Password</b> field.	Yes
JDBC Driver Path	Path to the JDBC driver that includes a JAR file. For example, /home/jars/db.jar on Linux.	Yes

### Advanced configuration parameters

<i>Table 36. Names and descriptions of the advanced configuration parameters</i>		
<b>Parameter name</b>	<b>Description</b>	<b>Mandatory field</b>
Java trace level	The trace levels that are used by the Java Custom providers. The default value is 2.	Yes

### Java API Client Configuration parameters

<i>Table 37. Names and descriptions of the Java API Client Configuration parameters</i>		
<b>Parameter name</b>	<b>Description</b>	<b>Mandatory field</b>
Class path for external jars	The path for jars required by Java API data provider that are not included with the agent.	No

## Configuring JBoss monitoring

The Monitoring Agent for JBoss offers a central point of management for your JBoss environment or application. The software provides a comprehensive means for gathering the information that is required to detect problems early and to prevent them. Information is standardized across the system. You can monitor multiple servers from a single console. By using the JBoss agent you can easily collect and analyze JBoss specific information.

### Before you begin

- Make sure that the system requirements for the JBoss agent are met in your environment. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\) for the JBoss agent](#).
- Before you configure the JBoss agent, the JBoss server first must be configured by completing the following tasks.
  1. [“Enable JMX MBean Server Connections” on page 318](#).
  2. [“Add a JBoss Server Management User” on page 317](#).



3. [“Enabling Web/HTTP Statistic Collection” on page 319](#) . This procedure is for JBoss EAP version 7.x and WildFly versions 8.x, 9.x and 10.x.

### About this task

The Managed System Name includes the instance name that you specify, for example, *instance\_name* : *host\_name* : *pc* , where *pc* is your two character product code. The Managed System Name is limited to 32 characters.

The instance name that you specify is limited to 28 characters minus the length of your host name. For example, if you specify JBoss as your instance name, your managed system name is JBoss:hostname:JE .

**Note:** If you specify a long instance name, the Managed System name is truncated and the agent code does not display correctly.

The JBoss agent is a multiple-instance agent. You must create an agent instance for each JBoss server you monitor, and start each agent instance manually.

### Procedure

1. Configure the agent on Windows systems by using the **IBM Performance Management** window or by using the silent response file.
  - [“Configuring the agent on Windows systems” on page 322](#) .
  - [“Configuring the agent by using the silent response file” on page 323](#) .
2. Configure the agent on Linux systems by running command line script and responding to prompts, or by using the silent response file.
  - [“Configuring the agent by responding to prompts” on page 322](#) .
  - [“Configuring the agent by using the silent response file” on page 323](#) .

### What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

If you are unable to view the data in the agent dashboards, first check the server connection logs and then the data provider logs. The default paths to these logs are as follows.

- **Linux** /opt/ibm/apm/agent/logs
- **Windows** C:\IBM\APM\TMAITM6\_x64\logs

## Add a JBoss Server Management User

Before the JBoss agent can gather data from the JBoss server, a management user must be added if one does not exist.

### Procedure

Use the JBoss **add-user** script to add a management user.

1. Go to the binary or bin directory under the JBoss server installation directory.
2. Run the **add-user** script.
  - **Linux** ./add-user.sh
  - **Windows** add-user.bat
3. Follow the prompts to generate a management user.

## Example

```
root@jboss-wf10-rh7:/apps/wildfly-10.0.0.Final/bin
] ./add-user.sh

What type of user do you wish to add?
  a) Management User (mgmt-users.properties)
  b) Application User (application-users.properties)
(a): a

Enter the details of the new user to add.
Using realm 'ManagementRealm' as discovered from the existing property files.
Username : MyAdmin
Password recommendations are listed below. To modify these restrictions edit the add-
user.properties
configuration file.
  - The password should be different from the username
  - The password should not be one of the following restricted values {root, admin,
administrator}
  - The password should contain at least 8 characters, 1 alphabetic character(s), 1 digit(s),
1 non-alphanumeric symbol(s)
Password :
Re-enter Password :
What groups do you want this user to belong to? (Please enter a comma separated list, or leave
blank
for none)[ ]:
About to add user 'MyAdmin' for realm 'ManagementRealm'
Is this correct yes/no? yes
Added user 'MyAdmin' to file '/apps/wildfly-10.0.0.Final/standalone/configuration/mgmt-
users.properties'
Added user 'MyAdmin' to file '/apps/wildfly-10.0.0.Final/domain/configuration/mgmt-
users.properties'
Added user 'MyAdmin' with groups to file
'/apps/wildfly-10.0.0.Final/standalone/configuration/mgmt-groups.properties'
Added user 'MyAdmin' with groups to file
'/apps/wildfly-10.0.0.Final/domain/configuration/mgmt-groups.properties'
Is this new user going to be used for one AS process to connect to another AS process?
e.g. for a slave host controller connecting to the master or for a Remoting connection for
server to
server EJB calls.
yes/no? no
```

## Enable JMX MBean Server Connections

Before the JBoss agent can gather data from the JBoss server, Java Management Extensions (JMX) MBean server connections must be enabled.

### Procedure

Follow the steps for your JBoss server release and version.

- Configure EAP 5.2.

Make a backup copy of the `run.conf` file, then add the following lines to it:

```
JAVA_OPTS="$JAVA_OPTS -Djboss.platform.mbeanserver"
JAVA_OPTS="$JAVA_OPTS -Dcom.sun.management.jmxremote"
JAVA_OPTS="$JAVA_OPTS -Dcom.sun.management.jmxremote.ssl=false"
JAVA_OPTS="$JAVA_OPTS -Dcom.sun.management.jmxremote.authenticate=false"
JAVA_OPTS="$JAVA_OPTS -Dcom.sun.management.jmxremote.port=1090"
JAVA_OPTS="$JAVA_OPTS -Djavax.management.builder.initial=
org.jboss.system.server.jmx.MBeanServerBuilderImpl"
```

- Configure AS 6.x.

Specify the bind address as a parameter when you start the JBoss server.

- **Linux** `jboss_server_home/bin/run.sh -b Ip_address`
- **Windows** `jboss_server_home\bin\run.bat -b <Ip_address>`

where `jboss_server_home` is the JBoss server installation directory.

For example, if the bind address is 10.77.9.250:

```
/apps/wildfly-9.0.2.Final/bin/run.sh -b 10.77.9.250
```

- Configure all other supported versions.

JBoss and WildFly servers are installed with their JMX ports disabled for remote management by default. You must change the configuration of the JBoss server to allow remote management. You must edit the `jboss_server_home/standalone/configuration/standalone.xml` to allow remote management.

- a) Make a backup copy of `jboss_server_home/standalone/configuration/standalone.xml` file.

Where `jboss_server_home` is the JBoss server installation directory.

- b) Allow remote configuration.

Search for `urn:jboss:domain:jmx` and within its subsystem section, make sure that the `remoting-connector` entry has `use-management-endpoint="true"`.

Example result.

```
<subsystem xmlns="urn:jboss:domain:jmx:1.3">
  <expose-resolved-model/>
  <expose-expression-model/>
  <remoting-connector use-management-endpoint="true"/>
</subsystem>
```

- c) Allow remote connections.

Find where the interfaces are defined and replace `127.0.0.1` (loopback) with the external IP on the server to bind to. Do not bind to `0.0.0.0`.

Example before replacement.

```
<interfaces>
  <interface name="management">
    <inet-address value="${jboss.bind.address.management:127.0.0.1}"/>
  </interface>
  <interface name="public">
    <inet-address value="${jboss.bind.address:127.0.0.1}"/>
  </interface>
  ...
```

Example after the replacement if the external IP address is `192.168.101.1`.

```
<interfaces>
  <interface name="management">
    <inet-address value="${jboss.bind.address.management:192.168.101.1}"/>
  </interface>
  <interface name="public">
    <inet-address value="${jboss.bind.address:192.168.101.1}"/>
  </interface>
  ...
```

## Enabling Web/HTTP Statistic Collection

Before the JBoss agent can gather JBoss server web metrics and other subsystem metrics, statistics collection must be enabled for each subsystem. This procedure is for JBoss EAP version 7.x and WildFly versions 8.x, 9.x and 10.x.

### Procedure

The **statistics-enabled** attribute of various JBoss subsystems controls statistic collection. This setting can be viewed and updated by using the JBoss command line interface.

**Note:** This procedure is for JBoss EAP version 7.x and WildFly versions 8.x, 9.x and 10.x.

1. Go to the binary or `bin` directory under the JBoss server installation directory.
2. Start the JBoss command line interface.

- `Linux` `./jboss-cli.sh --connect [--controller=IP:port]`

- **Windows** `jboss-cli.bat --connect [--controller=IP:port]`

where *IP* is the JBoss server's IP address and *port* is the JBoss server's port. For example, 192.168.10.20:9990.

**Tip:** If the connection attempt results in the error, "Failed to connect to the controller: The controller is not available at localhost:9990: java.net.ConnectException: WFLYPRT0053: Could not connect to http-remoting://localhost:9990. The connection failed: WFLYPRT0053: Could not connect to http-remoting://localhost:9990. The connection failed: Connection refused", use the **--controller** parameter.

This error indicates that the management server is not listening on the localhost IP address (127.0.0.1) and is configured to listen on the computer's IP address.

3. Run the following commands to view the current state of each subsystem's statistics-enabled attribute:

```
/subsystem=ejb3:read-attribute(name=enable-statistics)
/subsystem=transactions:read-attribute(name=statistics-enabled)
/subsystem=undertow:read-attribute(name=statistics-enabled)
/subsystem=webservices:read-attribute(name=statistics-enabled)
/subsystem=datasources/data-source=Data_Source_Name:read-attribute(name=statistics-enabled)
/subsystem=datasources/data-source=Data_Source_Name/statistics=pool:read-attribute(name=statistics-enabled)
/subsystem=datasources/data-source=Data_Source_Name/statistics=jdbc:read-attribute(name=statistics-enabled)
```

where *Data\_Source\_Name* is the name of a data source that is configured for use with JBoss.

**Note:** Data sources can be listed by using the command `/subsystem=datasources:read-resource`.

Example result when statistics are not enabled:

```
{
  "outcome" => "success",
  "result" => false
}
```

4. Run the following command to change the value of each subsystem's statistics-enabled attribute to *true*:

```
/subsystem=ejb3:write-attribute(name=enable-statistics, value=true)
/subsystem=transactions:write-attribute(name=statistics-enabled,value=true)
/subsystem=undertow:write-attribute(name=statistics-enabled,value=true)
/subsystem=webservices:write-attribute(name=statistics-enabled,value=true)
/subsystem=datasources/data-source=Data_Source_Name:write-attribute(name=statistics-enabled,value=true)
/subsystem=datasources/data-source=Data_Source_Name/statistics=pool:write-attribute(name=statistics-enabled,value=true)
/subsystem=datasources/data-source=Data_Source_Name/statistics=jdbc:write-attribute(name=statistics-enabled,value=true)
```

Example result when you enable statistics for a subsystem:

```
{
  "outcome" => "success",
```

```

    "response-headers" => {
      "operation-requires-reload" => true,
      "process-state" => "reload-required"
    }
  }
}

```

5. Exit the JBoss command line interface.
6. Restart the JBoss server.

**Note:** Any currently running JBoss agents with transaction tracking enabled must be restarted.

## Configuration Parameters for the JBoss agent

The configuration parameters for the JBoss agent are displayed in a table.

1. [JBoss Agent Settings](#) - JBoss agent environment settings.
2. [Table 39 on page 321](#) - Example JMX service URLs.

Parameter name	Description	Silent configuration file parameter name
Server Name	Provide a name to identify the JBoss/WildFly Server.	<b>KJE_SERVER</b>
Java home	The path to where Java is installed.	<b>JAVA_HOME</b>
JMX user ID	The user ID for connecting to the MBean server.	<b>KQZ_JMX_JSR160_JSR160_USER_ID</b>
JMX password	Password	<b>KQZ_JMX_JSR160_JSR160_PASSWORD</b>
JMX service URL	The service URL for connecting to the MBean server. See <a href="#">Table 39 on page 321</a> for examples.	<b>KQZ_JMX_JSR160_JSR160_SERVICE_URL</b>
JMX class path	The JAR files that are searched to locate a class or resource. Locate and enter the path to the <code>jboss-client.jar</code> file for your JBoss server. Example for a JBoss EAP 6 server, <code>/opt/EAP-6.3.0/jboss-eap-6.3/bin/client/jboss-client.jar</code> .	<b>KQZ_JMX_JSR160_JSR160_JAR_FILES</b>

JBoss server version	JMX service URL with default port <sup>1</sup>
WildFly 8, 9 and 10 JBoss EAP 7	service:jmx:remote+http://ip:9990 service:jmx:remote+https://ip:9994
JBoss EAP 6 JBoss AS 7	service:jmx:remoting-jmx://ip:9999
JBoss EAP 5.2 JBoss AS 6.1	service:jmx:rmi:///jndi/rmi://ip:1090/jmxrmi

<sup>1</sup> The port is based on the port in the JBoss configuration file entry `<socket-binding name="management-native" interface="management" port="$ {jboss.management.native.port:MMNN}"/>`. If the port was changed from the default value, adjust it according to the port number in your configuration file.

## Configuring the agent on Windows systems

You can configure the JBoss agent on Windows operating systems by using the **IBM Performance Management** window. After you update the configuration values, you must start the agent to save the updated values.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Cloud App Management**.
2. In the **IBM Performance Management** window, right-click the **Monitoring Agent for JBoss** template, and then click **Configure agent**.  
**Remember:** After you configure an agent instance for the first time, the **Configure agent** option is disabled. To configure the agent instance again, right-click on it and then click **Reconfigure...**
3. In the Monitoring Agent for JBoss window, complete the following steps:
  - a) Enter a unique instance name for the Monitoring Agent for JBoss instance, and click **OK**.
4. Enter the JBoss Server settings, then click **Next**.  
See [Table 38 on page 321](#) for an explanation of each of the configuration parameters.
5. Enter the Java settings, then click **Next**.  
See [Table 38 on page 321](#) for an explanation of each of the configuration parameters.
6. Enter the JMX settings, then click **Next**.  
See [Table 38 on page 321](#) for an explanation of each of the configuration parameters.
7. View the JBoss agent data collector settings.  
Leave the **DC Runtime Directory** blank during initial configuration of the agent. See [Table 38 on page 321](#) for an explanation of each of the configuration parameters.
8. Click **OK** to complete the agent configuration.
9. In the IBM Cloud App Management window, right-click the instance that you configured, and then click **Start**.

## Configuring the agent by responding to prompts

After installation of the JBoss agent, you must configure it before you start the agent. If the JBoss agent is installed on a local Linux or UNIX computer, you can follow these instructions to configure it interactively through command line prompts.

### About this task

**Remember:** If you are reconfiguring a configured agent instance, the value that is set in the last configuration is displayed for each setting. If you want to clear an existing value, press the space key when the setting is displayed.

### Procedure

1. On the command line, run the following command:

```
install_dir/bin/jboss-agent.sh config instance_name
```

where *install\_dir* is the path where the agent is installed and *instance\_name* is the name that you want to give to the agent instance.

Example

```
/opt/ibm/apm/agent/bin/jboss-agent.sh config example-inst01
```

2. Respond to the prompts to set configuration values for the agent.

See [“Configuration Parameters for the JBoss agent” on page 321](#) for an explanation of each of the configuration parameters.

3. Run the following command to start the agent:

```
install_dir/bin/jboss-agent.sh start instance_name
```

where *install\_dir* is the path where the agent is installed and *instance\_name* is the name of the agent instance.

Example

```
/opt/ibm/apm/agent/bin/jboss-agent.sh start example-inst01
```

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

The silent response file contains the agent configuration parameters with default values that are defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters.

After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

- To configure the JBoss agent in the silent mode, complete the following steps:
  - In a text editor, open the `jboss_silent_config.txt` file that is available at the following path:

- **Linux** | **UNIX** `install_dir/samples/jboss_silent_config.txt`
- **Windows** `install_dir\samples\jboss_silent_config.txt`

where *install\_dir* is the path where the agent is installed.

The default *install\_dir* paths are listed here:

- **Linux** `/opt/ibm/apm/agent`
- **Windows** `C:\IBM\APM\TMAITM6_x64`

Example

```
Linux | UNIX /opt/ibm/apm/agent/samples/jboss_silent_config.txt
```

```
Windows C:\IBM\APM\samples\jboss_silent_config.txt
```

- In the `jboss_silent_config.txt` file, specify values for all mandatory parameters. You can also modify the default values of other parameters.

See “[Configuration Parameters for the JBoss agent](#)” on page 321 for an explanation of each of the configuration parameters.

- Save and close the `jboss_silent_config.txt` file, and run the following command:

- **Linux** | **UNIX** `install_dir/bin/jboss-agent.sh config instance_name install_dir/samples/jboss_silent_config.txt`
- **Windows** `install_dir\bin\jboss-agent.bat config instance_name install_dir\samples\jboss_silent_config.txt`

where *install\_dir* is the path where the agent is installed and *instance\_name* is the name of the agent instance.

The default *install\_dir* paths are listed here:

- **Linux** /opt/ibm/apm/agent
- **Windows** C:\IBM\APM\TMAITM6\_x64

**Important:** Ensure that you include the absolute path to the silent response file. Otherwise, the agent data is not shown in the dashboards.

Example

```
Linux UNIX /opt/ibm/apm/agent/bin/jboss-agent.sh config example-inst01 /opt/ibm/apm/agent/samples/jboss_silent_config.txt
```

```
Windows C:\IBM\APM\bin\jboss-agent.bat config example-inst01 C:\IBM\APM\samples\jboss_silent_config.txt
```

d) Run the following command to start the agent:

- **Linux** **UNIX** *install\_dir*/bin/jboss-agent.sh start **instance\_name**
- **Windows** *install\_dir*\bin\jboss-agent.bat start **instance\_name**

where *install\_dir* is the path where the agent is installed and *instance\_name* is the name of the agent instance.

The default *install\_dir* paths are listed here:

- **Linux** /opt/ibm/apm/agent
- **Windows** C:\IBM\APM\TMAITM6\_x64

Example

```
Linux UNIX /opt/ibm/apm/agent/bin/jboss-agent.sh start example-inst01
```

```
Windows C:\IBM\APM\bin\jboss-agent.bat start example-inst01
```

## Configuring Linux KVM monitoring

You must configure the Monitoring Agent for Linux KVM to collect data of the Red Hat Enterprise Virtualization Hypervisor (RHEVH) and Red Hat Enterprise Virtualization Manager (RHEVM) servers. After you install the agent on a server or a virtual machine, you must create the first instance, and start the agent manually.

### Before you begin

Review hardware and software prerequisites. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\)](#) for the Monitoring Agent for Linux KVM.

### About this task

The Linux KVM agent is a multi-instance and multi-connection agent. Multi-instance means that you can create multiple instances and each instance can make multiple connections to one or more RHEVM or RHEVH servers.

**Remember:** Use different instances to monitor RHEVM or the RHEVH servers.

You can use the same configuration script to configure instances for the RHEVH and the RHEVM servers:

- To configure a connection to the RHEVM server, complete the steps that are mentioned in the "Configuring a connection to the RHEVM server" topic.
- To configure a connection to the RHEVH server, complete the steps that are mentioned in the "Configuring a connection to the RHEVH server" topic.



## Creating a user and granting required permissions

Before you configure the Linux KVM agent, you must create a user and grant required permissions to the user to monitor the RHEVM and RHEVH servers.

### Procedure

1. Open the **Red Hat Enterprise Virtualization Manager Web Administration** portal.
2. Click **Configure**.
3. In the **Configuration** window, select **Roles**.
  - a) To create a role, click **New**.
  - b) In the **New Role** window, add the name of the role and select **Admin** as the account type.
  - c) Ensure that the check boxes in the **Check boxes to Allow Action** pane are not selected, and click **OK**.
4. In the **Configuration** window, select **System Permission**.
  - a) To grant a user permission, click **Add**.
  - b) In the **Add System Permission to User** window, select the user to whom you want to grant the permission.
  - c) From the **Assign role to user** list, select the role that you created and click **OK**.

### What to do next

Complete the agent configuration:

- [“Configuring a connection to the RHEVH server” on page 331](#)
- [“Configuring a connection to the RHEVM server” on page 329](#)

## Configuring protocols

The agent uses different protocols to connect to the RHEVH server. You can configure any of these protocols: SSH, TLS, or TCP.

### About this task

The Linux KVM agent remotely connects to each hypervisor by using the **virsh** tool that manages your QEMU-KVM virtual machines, and collects metrics. The `libvirt` API in the agent environment uses several different remote transport protocols. For the list of supported protocols, see the [Remote support page](#).

### Configuring the SSH protocol

You can configure the SSH protocol to remotely monitor a host.

### About this task

**Assumption:** The Linux KVM agent is installed on host A. You want to remotely monitor the hypervisor on host B.

### Procedure

1. Log in to host A with the same user ID that runs the Linux KVM agent process, for example, the root user ID.

**Tip:** Ensure that you know the ID on host B that accepts the SSH connection and the root user ID on host A.
2. Generate the **id\_rsa** and **id\_rsa.pub** keys on host A by using the `ssh-keygen` utility.

The keys are saved at the following location: `~/.ssh: $ ssh-keygen -t rsa`.
3. Copy the authorized keys from host B:

```
$ scp Id on host B@name or IP address of host B:~/.ssh/authorized_keys
~/.ssh/authorized_keys_from_B
```

4. Append the public key for host A to the end of the authorized keys for host B:

```
cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys_from_B
```

5. Copy the authorized keys back to host B:

```
$ scp ~/.ssh/authorized_keys_from_B Id on host B@name or IP address of host
B:~/.ssh/authorized_keys
```

**Remember:** If you are monitoring multiple hosts, repeat steps “3” on page 325, “4” on page 326, and “5” on page 326 for each host.

6. Remove the authorized keys that you copied on host B:

```
~/.ssh/authorized_keys_from_B
```

7. Add the following command to the `~/.bash_` profile of the current ID on host A:

```
$ eval `ssh-agent`
```

**Remember:** Ensure that you use the single back quotation mark ( ` ) that is located under the tilde ( ~ ) on US keyboards, rather than the single quotation mark ( ' ).

8. Add the identity to host A and enter the password that you used when you created the ID:

```
$ ssh-add ~/.ssh/id_rsa
```

9. Run the following command if you receive the Could not open a connection to your authentication agent message:

```
exec ssh-agent bash
```

**Tip:** You can replace the bash with the shell that you are using and then run the following command again:

```
$ ssh-add ~/.ssh/id_rsa
```

10. Test the SSH protocol to ensure that it connects from host A to host B without entering the SSH password:

**Tip:** If you are monitoring multiple hosts, use the following command to test the connection for each host:

```
$ ssh Id on host B@name or IP address of host B
```

11. To verify the connection, run the following command:

```
virsh -c qemu+ssh://Id on host B@name or IP address of host B:port/system
```

If you did not change the default SSH port, omit the `:port` section of the command.

**Important:** If the `virsh` command succeeds, the Linux KVM agent connects to the hypervisor.

12. You must restart host A before you restart the Linux KVM agent on host A. To restart, run the `ssh-add` command again and specify the password each time.

**Tip:** You can use SSH keychains to avoid reentering the password.

### Configuring the TLS protocol

You can configure the TLS protocol to remotely monitor a host.

#### About this task

**Assumption:** The Linux KVM agent is installed on host A. You want to remotely monitor the hypervisor on host B.

#### Procedure

1. To create a certificate authority (CA) key and a certificate in your hypervisor, complete the following steps:

- a) Log in to host B.
  - b) Create a temporary directory and change the path to this temporary directory:
 

```
mkdir cert_files
cd cert_files
```
  - c) Create a 2048-bit RSA key:
 

```
openssl genrsa -out cakey.pem 2048
```
  - d) Create a self-signed certificate to your local CA:
 

```
openssl req -new -x509 -days 1095 -key cakey.pem -out \
cacert.pem -sha256 -subj "/C=US/L=Austin/O=IBM/CN=my CA"
```
  - e) Check your CA certificate:
 

```
openssl x509 -noout -text -in cacert.pem
```
2. To create the client and server keys and certificates in your hypervisor, complete the following steps:
- a) Create the keys:
 

```
openssl genrsa -out serverkey.pem 2048
openssl genrsa -out clientkey.pem 2048
```
  - b) Create a certificate signing request for the server:
 

**Remember:** Change the kvmhost.company.org address, which is used in the server certificate request, to the fully qualified domain name of your hypervisor host.

```
openssl req -new -key serverkey.pem -out serverkey.csr \
-subj "/C=US/O=IBM/CN=kvmhost.company.org"
```
  - c) Create a certificate signing request for the client:
 

```
openssl req -new -key clientkey.pem -out clientkey.csr \
-subj "/C=US/O=IBM/OU=virtualization/CN=root"
```
  - d) Create client and server certificates:
 

```
openssl x509 -req -days 365 -in clientkey.csr -CA cacert.pem \
-CAkey cakey.pem -set_serial 1 -out clientcert.pem
openssl x509 -req -days 365 -in serverkey.csr -CA cacert.pem \
-CAkey cakey.pem -set_serial 94345 -out servercert.pem
```
  - e) Check the keys:
 

```
openssl rsa -noout -text -in clientkey.pem
openssl rsa -noout -text -in serverkey.pem
```
  - f) Check the certificates:
 

```
openssl x509 -noout -text -in clientcert.pem
openssl x509 -noout -text -in servercert.pem
```
3. To distribute the keys and certificates to the host server, complete the following steps:
- a) Copy the CA certificate cacert.pem file to this directory: /etc/pki/CA
 

```
cp cacert.pem /etc/pki/CA/cacert.pem
```
  - b) Create the /etc/pki/libvirt directory, and copy the servercert.pem server certificate file to the /etc/pki/libvirt directory. Ensure that only the root user can access the private key.
 

```
mkdir /etc/pki/libvirt
cp servercert.pem /etc/pki/libvirt/
chmod -R o-rwx /etc/pki/libvirt
```

**Remember:** If the keys or certificates are named incorrectly or copied to the wrong directories, the authorization fails.

- c) Create the `/etc/pki/libvirt/private` directory and copy the `serverkey.pem` server key file to the `/etc/pki/libvirt/private` directory. Ensure that only the root user can access the private key.

```
mkdir /etc/pki/libvirt/private
```

```
cp serverkey.pem /etc/pki/libvirt/private/.
```

```
chmod -R o-rwx /etc/pki/libvirt/private
```

**Remember:** If the keys or certificates are named incorrectly or copied to the wrong directories, the authorization fails.

- d) Verify that the files are correctly placed:

```
find /etc/pki/CA/*|xargs ls -l
```

```
ls -lR /etc/pki/libvirt
```

```
ls -lR /etc/pki/libvirt/private
```

**Remember:** If the keys or certificates are named incorrectly or copied to the wrong directories, the authorization fails.

4. To distribute keys and certificates to clients or management stations, complete the following steps:

- a) Log in to host A.

- b) Copy the CA certificate `cacert.pem` from the host to the `/etc/pki/CA` directory in host A without changing the file name.

```
scp kvmhost.company.org:/tmp/cacert.pem /etc/pki/CA/
```

- c) Copy the client certificate `clientcert.pem` file to the `/etc/pki/libvirt` directory from host B. Use the default file names and make sure that only the root user is able to access the private key.

```
mkdir /etc/pki/libvirt/
```

```
scp kvmhost.company.org:/tmp/clientcert.pem /etc/pki/libvirt/.
```

```
chmod -R o-rwx /etc/pki/libvirt
```

**Remember:** If the keys or certificates are named incorrectly or copied to the wrong directories, the authorization fails.

- d) Copy the client key `clientkey.pem` to the `/etc/pki/libvirt/private` directory from the host. Use the default file names and ensure that only the root user can access the private key.

```
mkdir /etc/pki/libvirt/private
```

```
scp kvmhost.company.org:/tmp/clientkey.pem /etc/pki/libvirt/private/.
```

```
chmod -R o-rwx /etc/pki/libvirt/private
```

**Remember:** If the keys or certificates are named incorrectly or copied to the wrong directories, the authorization fails.

- e) Verify that the files are correctly placed:

```
ls -lR /etc/pki/libvirt
```

```
ls -lR /etc/pki/libvirt/private
```

5. To edit the `libvirtd` daemon configuration, complete the following steps:

- a) Log in to host B.

- b) Make a copy of the `/etc/sysconfig/libvirtd` file and the `/etc/libvirt/libvirtd.conf` file.

- c) Edit the `/etc/sysconfig/libvirtd` file and ensure that the **--listen** parameter is passed to the `libvirtd` daemon. This step ensures that the `libvirtd` daemon is listening to network connections.
  - d) Edit the `/etc/libvirt/libvirtd.conf` file and configure a set of allowed subjects with the **tls\_allowed\_dn\_list** directive in the `libvirtd.conf` file.
 

**Important:** The fields in the subject must be in the same order that you used to create the certificate.
  - e) Restart the `libvirtd` daemon service for changes to take effect:
 

```
/etc/init.d/libvirtd restart
```
6. To change the firewall configuration, access the security level configuration and add TCP port 16514 as a trusted port.
  7. To verify that the remote management is working, run the following command on host A:
 

```
virsh -c qemu+tls://kvmhost.company.org/system list --all
```

### Configuring the TCP protocol

Use the TCP protocol only for testing.

#### About this task

**Assumption:** The Linux KVM agent is installed on host A. You want to remotely monitor the hypervisor on host B.

#### Procedure

1. Log in to host B.
2. Edit the `/etc/libvirt/libvirtd.conf` file and ensure that the **listen\_tcp** parameter is enabled, and the value of the **tcp\_port** parameter is set to the default value of 16509.
3. Edit the `/etc/libvirt/libvirtd.conf` file to set the **auth\_tcp** parameter to "none". This step instructs TCP not to authenticate the connection.
4. Restart the **libvirt** daemon on host B in listening mode by running it with the **--listen** flag or by editing the `/etc/sysconfig/libvirtd` file and uncommenting the `LIBVIRT_ARGS="--listen"` line.
5. To verify the connection, run the following command:

```
virsh -c qemu+tcp://kvmhost.company.org:port/system
```

If you did not change the default TCP port, omit the **:port** section of the command.

**Important:** If the **virsh** command succeeds, the Linux KVM agent connects to the hypervisor.

#### What to do next

Configure the agent by completing the steps that are described in [“Configuring a connection to the RHEVH server”](#) on page 331.

## Configuring a connection to the RHEVM server

To configure a connection to the RHEVM server, you must run the script and respond to prompts.

#### Before you begin

1. Download the security certificate that is available at the following path:

```
https://RHEVM-HOST:RHEVM-PORT/ca.crt
```

Where

#### RHEVM-HOST

The name of the host.

**RHEVM-PORT**

The port that you use in your RHEVM environment.

2. Use the *keytool* utility to import the security certificate file to generate a local keystore file:

```
keytool -import -alias ALIAS -file CERTIFICATE_FILE -keystore  
KEYSTORE_FILE
```

Example **keytool -import -alias RHEVM36vmwt9 -file hjs495-vmw-t-9.cer  
-keystore RHEVM36KeyStore**

Where

**ALIAS**

A unique reference for each certificate that is added to the certificate truststore of the agent, for example, an appropriate alias for the certificate from *datasource.example.com* is *datasource*.

**CERTIFICATE\_FILE**

The complete path and file name to the data source certificate that is being added to the truststore.

**KEYSTORE\_FILE**

The name of the keystore file that you want to specify.

**Tip:** The *keytool* utility is available with Java Runtime Environment (JRE). The keystore file is stored at the same location from where you run the command.

3. Ensure that the user, who connects to the RHEVM, is an administrator with the SuperUser role. Use can use an existing user ID with this role, or you can create a new user ID by completing the steps that are mentioned in [“Creating a user and granting required permissions”](#) on page 325.

**Procedure**

1. On the command line, run the following command:

```
install_dir/bin/linux_kvm-agent.sh config instance_name
```

Example **/opt/ibm/apm/agent/bin/linux\_kvm-agent.sh config instance\_name**

Where

**instance\_name**

The name that you want to give to the instance.

**install\_dir**

The path where the agent is installed.

2. Respond to the prompts and specify values for the configuration parameters.

For information about the configuration parameters, see [“Configuration parameters to connect to the RHEVM server”](#) on page 331.

3. Run the following command to start the agent:

```
install_dir/bin/linux_kvm-agent.sh start instance_name
```

Example **/opt/ibm/apm/agent/bin/linux\_kvm-agent.sh start instance\_name**

**What to do next**

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

If you need help with troubleshooting, see the [IBM Cloud App Management & IBM Cloud Application Performance Management](#) on developerWorks®.

## Configuring a connection to the RHEVH server

To configure a connection to the RHEVH server, you must run the script and respond to prompts.

### Before you begin

- Ensure that the user, who connects to the RHEVM, is a root user. You can use an existing user ID or create a new user ID by completing the steps that are mentioned in [“Creating a user and granting required permissions”](#) on page 325.
- Configure the protocol that you want to use to connect to the RHEVH server by completing the steps that are described in [“Configuring protocols”](#) on page 325.

### Procedure

1. On the command line, run the following command:

```
install_dir/bin/linux_kvm-agent.sh config instance_name
```

Example `/opt/ibm/apm/agent/bin/linux_kvm-agent.sh config instance_name`

Where

#### **instance\_name**

The name that you want to give to the instance.

#### **install\_dir**

The path where the agent is installed.

2. Respond to the prompts and specify values for the configuration parameters.

For information about the configuration parameters, see [“Configuration parameters to connect to the RHEVH server”](#) on page 333.

3. Run the following command to start the agent:

```
install_dir/bin/linux_kvm-agent.sh start instance_name
```

Example `/opt/ibm/apm/agent/bin/linux_kvm-agent.sh start instance_name`

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

If you need help with troubleshooting, see the [IBM Cloud App Management & IBM Cloud Application Performance Management](#) on developerWorks.

## Configuration parameters to connect to the RHEVM server

You can modify the default values of configuration parameters that are used for connecting the agent with the RHEVM server.

The following table contains detailed descriptions of the configuration parameters.

<b>Parameter name</b>	<b>Description</b>	<b>Mandatory field</b>
Edit Monitoring Agent for Linux KVM settings	Indicates that you can begin editing the default values of the configuration parameters. Enter 1 (Yes), which is also the default value, to continue.	Yes
Maximum number of Data Provider Log Files	The maximum number of log files that the data provider creates before it overwrites the previous log files. The default value is 10.	Yes

Table 40. Names and descriptions of the configuration parameters for connecting to the RHEVM server (continued)

Parameter name	Description	Mandatory field
Maximum Size in KB of Each Data Provider Log	The maximum size in KB that a data provider log file must reach before the data provider creates a new log file. The default value is 5190 KB.	Yes
Level of Detail in Data Provider Log	<p>The level of detail that can be included in the log file that the data provider creates. The default value is 4 (Info). The following values are valid:</p> <ul style="list-style-type: none"> <li>• 1 = Off: No messages are logged.</li> <li>• 2 = Severe: Only errors are logged.</li> <li>• 3 = Warning: All errors and messages that are logged at the Severe level and potential errors that might result in undesirable behavior.</li> <li>• 4 = Info: All errors and messages that are logged at the Warning level and high-level informational messages that describe the state of the data provider when it is processed.</li> <li>• 5 = Fine: All errors and messages that are logged at the Info level and low-level informational messages that describe the state of the data provider when it is processed.</li> <li>• 6 = Finer: All errors and messages that are logged at the Fine level plus highly-detailed informational messages, such as performance profiling information and debug data. Selecting this option can adversely affect the performance of the monitoring agent. This setting is intended only as a tool for problem determination along with IBM support staff.</li> <li>• 7 = Finest: All errors and messages that are logged at the Fine level and the most detailed informational messages that include low-level programming messages and data. Choosing this option might adversely affect the performance of the monitoring agent. This setting is intended only as a tool for problem determination along with IBM support staff.</li> <li>• 8 = All: All errors and messages are logged.</li> </ul>	Yes
Edit Hypervisor settings	Indicates whether you want to edit the parameters for a connection to the RHEVH server. Enter 5 (Next) because you are configuring a connection to the RHEVM server. The default value is 5 (Next).	Yes
Edit RHEVM Connection Details settings	<p>Indicates whether you want to edit the parameters for a connection to the RHEVM server. Enter 1 (Add) to continue. The default value is 5 (Next).</p> <p><b>Important:</b> After you specify values for all the configuration parameters, you are again prompted to indicate whether you want to continue to edit the parameters. Enter 5 (Exit).</p>	Yes
RHEVM ID	The unique user name, which you specify for the RHEVM that you connect to.	Yes
Host	The host name or IP address of the data source that is used to connect to the RHEVM server.	Yes



*Table 40. Names and descriptions of the configuration parameters for connecting to the RHEVM server (continued)*

<b>Parameter name</b>	<b>Description</b>	<b>Mandatory field</b>
User	The user name of the data source with sufficient privileges to connect to the RHEVM server.	Yes
Password	The password of the user name that you use to connect to the RHEVM server.	Yes
Re-type password	The same password that you specified in the <b>Password</b> field.	Yes
Port	The port number that is used to connect to the RHEVM server.	Yes
Domain	The domain to which the user belongs.	Yes
KeyStorePath	The file path and name of the local keystore file that you created by using the <b>keytool</b> command.	Yes

### **Configuration parameters to connect to the RHEVH server**

You can modify the default values of configuration parameters that are used for connecting the agent with the RHEVH server.

The following table contains detailed descriptions of the configuration parameters.

*Table 41. Names and descriptions of the configuration parameters for connecting to the hypervisor*

<b>Parameter name</b>	<b>Description</b>	<b>Mandatory field</b>
Edit Monitoring Agent for Linux KVM settings	Indicates that you can begin editing the default values of the configuration parameters. Enter 1 (Yes), which is also the default value, to continue.	Yes
Maximum number of Data Provider Log Files	The maximum number of log files that the data provider creates before it overwrites the previous log files. The default value is 10.	Yes
Maximum Size in KB of Each Data Provider Log	The maximum size in KB that a data provider log file must reach before the data provider creates a new log file. The default value is 5190 KB.	Yes

Table 41. Names and descriptions of the configuration parameters for connecting to the hypervisor (continued)

Parameter name	Description	Mandatory field
Level of Detail in Data Provider Log	<p>The level of detail that can be included in the log file that the data provider creates. The default value is 4 (Info). The following values are valid:</p> <ul style="list-style-type: none"> <li>• 1 = Off: No messages are logged.</li> <li>• 2 = Severe: Only errors are logged.</li> <li>• 3 = Warning: All errors and messages that are logged at the Severe level and potential errors that might result in undesirable behavior.</li> <li>• 4 = Info: All errors and messages that are logged at the Warning level and high-level informational messages that describe the state of the data provider when it is processed.</li> <li>• 5 = Fine: All errors and messages that are logged at the Info level and low-level informational messages that describe the state of the data provider when it is processed.</li> <li>• 6 = Finer: All errors and messages that are logged at the Fine level plus highly-detailed informational messages, such as performance profiling information and debug data. Selecting this option can adversely affect the performance of the monitoring agent. This setting is intended only as a tool for problem determination along with IBM support staff.</li> <li>• 7 = Finest: All errors and messages that are logged at the Fine level and the most detailed informational messages that include low-level programming messages and data. Selecting this option might adversely affect the performance of the monitoring agent. This setting is intended only as a tool for problem determination along with IBM support staff.</li> <li>• 8 = All: All errors and messages are logged.</li> </ul>	Yes
Edit Hypervisor settings	Indicates whether you want to edit the parameters for a Hypervisor connection. Enter 1 (Add). The default value is 5 (Next).	Yes
Hypervisor ID	The unique user name, which you specify for the RHEVH that you connect to.	Yes
Host	The host name or IP address of the data source that is used to connect to the RHEVH server.	Yes
User	A user name of the data source with sufficient privileges to connect to the RHEVM server.	Yes
Remote Transport	<p>The protocol that is used by the local libvirt API to connect to remote libvirt APIs. The default value is 1. The following values are valid:</p> <ul style="list-style-type: none"> <li>• 1 = SSH</li> <li>• 2 = TLS</li> <li>• 3 = TCP (Unencrypted - not recommended for production use)</li> </ul>	Yes

*Table 41. Names and descriptions of the configuration parameters for connecting to the hypervisor (continued)*

Parameter name	Description	Mandatory field
Port	The port that is used by the transport protocol to connect to the libvirt API. The default value is 22.  <b>Important:</b> This port is only needed if the standard ports were changed (22 for SSH, 16514 for TLS, 16509 for TCP).	Yes
Domain	The domain to which the user belongs.	Yes
Connection Instance Type	Indicates whether the local libvirt API connects to the privileged system driver or the per-user unprivileged session driver. The default value is 1. The following values are valid: <ul style="list-style-type: none"> <li>• 1 = system</li> <li>• 2 = session</li> </ul>	Yes
Edit RHEVM Connection Details settings	Indicates whether you want to edit the parameters for a connection to the RHEVM server. Enter 1 (Add) to continue. The default value is 5 (Next).  <b>Important:</b> After you specify values for all the configuration parameters, you are again prompted to indicate whether you want to continue to edit the parameters. Enter 5 (Next).	Yes

## Configuring MariaDB monitoring

You must configure the MariaDB agent so that the agent can collect data to monitor the availability and performance of MariaDB server resources.

### Before you begin

- Ensure that the system requirements for the MariaDB agent are met in your environment.
- Ensure that a user is created in MariaDB database to run the agent. The user does not require any specific privileges on the MariaDB database that it monitors.

### About this task

The MariaDB agent is a single instance agent. You must configure the agent manually after it is installed. You can configure the agent on Windows and Linux operating systems. The agent requires an instance name and the MariaDB server user credentials to configure it. The instance name that you specify can contain up to 28 characters.

## Configuring the agent on Windows systems

You can configure the agent on Windows operating systems by using the IBM Cloud Application Performance Management window. After you update the configuration values, start the agent to apply the updated values.

### Procedure

To configure the agent on Windows operating systems, complete the following steps:

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, complete these steps:
  - a) Double-click the **Monitoring Agent for MariaDB** template.
  - b) In the **Monitoring Agent for MariaDB** window, specify an instance name and click **OK**.

3. In the **Monitoring Agent for MariaDB** window, complete these steps:
  - a) In the **IP Address** field, enter the IP address of MariaDB server that you want to monitor remotely.  
If the agent is installed on a server to be monitored, retain the default value.
  - b) In the **JDBC user name** field, enter the name of a MariaDB server user. The default value is root.
  - c) In the **JDBC password** field, type the password of a JDBC user.
  - d) In the **Confirm JDBC password** field, type the password again.
  - e) In the **JDBC Jar File** field, click **Browse** and locate the directory that contains the MariaDB connector Java file and select it.
  - f) Click **Next**.
  - g) In the **JDBC port number** field, specify the port number of the JDBC server.  
The default port number is 3306.
  - h) From the **Java trace level** list, select a trace level for Java.  
The default value is Error.
  - i) Click **OK**.  
The instance is displayed in the **IBM Performance Management** window.

4. Right-click the **Monitoring Agent for MariaDB** instance, and click **Start**.

**Remember:** To configure the agent again, complete these steps in the **IBM Performance Management** window:

- a. Stop the agent instance that you want to configure.
- b. Right-click the **Monitoring Agent for MariaDB** instance, and click **Reconfigure**.
- c. Repeat steps 3 and 4.

### What to do next

- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent on Linux systems

You can run the configuration script and respond to prompts to configure the agent on Linux operating systems.

### Procedure

To configure the agent on Linux operating systems, complete the following steps:

1. On command line, run the following command:

```
install_dir/bin/mariadb-agent.sh config instance_name
```

Where *instance\_name* is the name you want to give to the instance, and *install\_dir* is the installation directory for the MariaDB agent.

2. When you are prompted to enter a value for the following parameters, press **Enter** to accept the default value, or specify a different value and press **Enter**.
  - IP address
  - JDBC username
  - JDBC password
  - Retype JDBC password
  - JDBC JAR file
  - JDBC port number (Default port number is 3306.)

- Java trace level (Default value is Error.)

For more information about the configuration parameters, see [“Configuring the agent by using the silent response file”](#) on page 407.

3. Run the following command to start the agent:

```
install_dir/bin/mariadb-agent.sh start instance_name
```

### What to do next

- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

You can use the silent response file to configure the MariaDB agent on Linux and Windows system. After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

To configure the agent by using the silent response file, complete the following steps:

**Remember:** This procedure assumes the following default path where the agent is installed:

**Windows** C:\IBM\APM

**Linux** opt/ibm/apm/agent

If the agent is installed at a different path, substitute the path in the instructions. Also, edit the **AGENT\_HOME** parameter in the silent response file to specify the path where the agent is installed.

1. In a text editor, open the response file that is available at the following path:

**Linux** `install_dir/samples/mariadb_silent_config.txt`

**Windows** `install_dir\samples\mariadb_silent_config.txt`

Where `install_dir` is the installation directory of the MariaDB agent

2. In the response file, specify a value for the following parameters:

- For the **Server Name** parameter, specify the IP address of a MariaDB server that you want to monitor remotely. Otherwise, retain the default value as localhost.
- For the **JDBC user name** parameter, retain the default username value of root or specify the name of a user with privileges to view the INFORMATION\_SCHEMA tables.
- For the **JDBC password** parameter, enter the JDBC user password.
- For the **JDBC Jar File** parameter, retain the default path if this path to the MariaDB connector for the Java JAR file is correct. Otherwise, enter the correct path. The connector is available at the following default path:

**Linux** `/usr/share/java/mariadb-connector-java.jar`

**Windows** `C:\Program Files (x86)\MariaDB\mariadb-connector-java.jar`

- For the **JDBC port number** parameter, retain the default port number of 3306 or specify a different port number.
- For the **Java trace level** parameter, retain the default value of Error or specify a different level according to the IBM support instructions.

3. Save and close the response file, and run the following command to update the agent configuration settings:

```
Linux install_dir/bin/mariadb-agent.sh config instance_name install_dir/samples/mariadb_silent_config.txt
```

```
Windows install_dir\BIN\mariadb-agent.bat config instance_name install_dir\samples\mariadb_silent_config.txt
```

Where *instance\_name* is the name that you want to give to the instance, and *install\_dir* is the installation directory of MariaDB agent.

**Important:** Ensure to include the absolute path to the silent response file. Otherwise, dashboards do not display agent data.

### What to do next

- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring Microsoft .NET monitoring

When you install the Monitoring Agent for Microsoft .NET, the agent is automatically configured and starts with the default configuration settings.

### Before you begin

- Review the hardware and software prerequisites. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\)](#) for the Microsoft .NET agent.
- Ensure that the user, who connects to the Microsoft .NET environment or application, has administrator privileges. Use an existing user with administrator privileges, or create a new user. Assign administrator privileges to the new user by adding the new user to the Administrators group.

### About this task

The agent starts automatically after installation to collect the resource monitoring data. However, to configure the Microsoft .NET agent, you can use a local or a domain user provided that the user has administrator privileges. You can configure the agent on Windows operating systems.

The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history”](#) on page 52.

## Permissions to run an agent by using a local or domain account

Only a local or domain user who is a member of Administrators group has permissions to run the Microsoft .NET agent. This topic provides conditions that must be met if the local or domain user is not a member of Administrators group.

### User must have the following permissions to the system drive and agent installation drive

- Read
- Write
- Execute
- Modify

### User must have the following permission to the HKEY\_LOCAL\_MACHINE registry key

- Read

### User must be a member of the following groups on the monitored server

- Users
- IIS\_IUSRS
- Performance Monitor Users
- Performance Log Users

**Note:** However, it is advisable to run Microsoft .NET agent with a local or domain user that is a member of local Administrators group.

## Configuring the agent on Windows systems

You can configure the Microsoft .NET agent on Windows operating systems by using the **IBM Performance Management** window. After you update the configuration values, you must restart the agent to save the updated values.

### About this task

You can configure the agent when the agent is running or stopped. The agent remains in the same state after configuration. For example, if the agent is running, it remains in the running state after configuration.

The Microsoft .NET agent provides default values for some parameters. You can specify different values for these parameters.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Microsoft .NET agent**, and click **Reconfigure**.
3. In the **Restart of Monitoring Agent for Microsoft .NET** window, click **Yes**.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## Changing the user account

After you configure the Microsoft .NET agent, you can change the user account from the local user to the domain user.

### About this task

By default, the Microsoft .NET agent runs under the local user account.

### Procedure

1. Run the following command to verify which user ID is being used for starting the agent:

```
install_dir\InstallITM\KinCinfo.exe -r
```

2. If the monitoring agent was started with a user ID that does not belong to the Administrator group, stop the agent.
3. Open the **Manage Monitoring Services** window.
4. Right-click the agent instance, and click **Change Startup**.
5. Specify the fully qualified user ID as <Domain\User ID>, and then specify the password.
6. Start the Microsoft .NET agent.

## Configuring Microsoft Active Directory monitoring

The Monitoring Agent for Microsoft Active Directory is automatically configured and started after installation.

### Before you begin

- The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history” on page 52](#).

To view data for all attributes in the dashboard, complete the following tasks:

- [“Running the Microsoft Active Directory agent as an administrator user” on page 340](#)
- [“Configuring local environment variables” on page 341](#)

### About this task

You may choose to run the agent as non-administrator user, configure domain services for attribute group AD\_Services\_Status or upgrade the agent.

## Running the Microsoft Active Directory agent as an administrator user

You must have administrative rights to run the Microsoft Active Directory agent.

### About this task

All data sets are available to the users who are members of the Administrators group. In this task, you create a user, assign administrator rights to the user, and change the user account for the agent to this user.

### Procedure

1. Click **Start > All Programs > Administrative Tools > Active Directory Users and Computers**.
2. To expand the domain where you want to create the user, click the plus sign (+) next to the name of a domain.
3. Right-click **Users**, and then click **New > User**.
4. To create a new user, open the **New Object - User** wizard.  
By default, a new user is a member of the Domain Users group.
5. Right-click the new user that is created in the Domain Users group, and click **Properties**. The **Username Properties** window is displayed. The username is the name of the new user.
6. In the **Username Properties** window, complete the following steps:
  - a) Click the **Member of** tab. In the **Member of** area, add the Administrators group.
  - b) Click **Apply**, and then click **OK**.
7. Click **Start > Run**, and then type `services.msc`.
8. In the **Services** window, complete the following steps:
  - a) Right-click the **Monitoring Agent for Active Directory service**, and click **Properties**.
  - b) In the **Monitoring Agent for Active Directory Properties** window, on the **Log On** tab, click **This Account**. Enter the user credentials.
  - c) Click **Apply**, and then click **OK**.
9. Restart the agent service.



## Configuring local environment variables

You must specify values for the environment variables to view the Sysvol replication data in the dashboard. Optionally, you can also update the cache interval value to enable or disable caching.

### Procedure

1. In the **IBM Performance Management** window, from the **Actions** menu, click **Advanced > Edit ENV File**.
2. In the K3ZENV file, change the values of the following environment variables.

#### **ADO\_CACHE\_INTERVAL**

Determines whether to start or stop the caching and is used to set a value for the cache interval. Cache interval is the duration in seconds between two consecutive data collections. You can specify any positive integer value for the cache interval to start the caching. You can specify the zero value for the cache interval to stop the caching. By default, the caching is started, and the cache interval value is set to 1200.

#### **ADO\_SYSVOL\_FORCE\_REPLICATION\_FLAG**

Determines whether the force replication that is initiated by the agent is enabled or disabled. The default value of this variable is TRUE. To disable force replication, change the value of this variable to FALSE.

#### **ADO\_SYSVOL\_REPLICATION\_TEST\_INTERVAL**

Determines the time interval in minutes between two Sysvol replication tests. The default value of this variable is 0 minutes. To complete the Sysvol replication test, ensure that the value of this variable is greater than zero.

#### **ADO\_SYSVOL\_REPLICATION\_TEST\_VERIFICATION\_INTERVAL**

Determines the amount of time in minutes that the agent waits to verify the results of Sysvol replication after it completes the Sysvol replication test.

The value of the **ADO\_SYSVOL\_REPLICATION\_TEST\_INTERVAL** variable must be greater than the value of the **ADO\_SYSVOL\_REPLICATION\_TEST\_VERIFICATION\_INTERVAL** variable. You can use the following values for these variables:

**ADO\_SYSVOL\_REPLICATION\_TEST\_INTERVAL: 1440**

**ADO\_SYSVOL\_REPLICATION\_TEST\_VERIFICATION\_INTERVAL: 30**

After you assign valid values to the two environment variables, the Active Directory agent creates one file in the Sysvol shared folder of the managed system and initializes forced Sysvol replication. This forced replication is initialized from the managed system to the Sysvol shared folders of the Sysvol replication partners. After you verify the results of the replication test, the agent removes the files that are created and replicated from the managed system and Sysvol replication partners.

3. Optional: In the K3ZENV file, add the **APM\_ATTRIBUTES\_ENABLE\_COLLECTION** environmental variable and set its value to Yes to view data for the following data sets in the **Attribute details** tab.
  - Services
  - Replication
  - File Replication Service
  - Moved or Deleted Org Unit
  - LDAP
  - Security Accounts Manager
  - DFS
  - Address Book
  - Event Log
  - Password Setting Objects

**Remember:** If you want to disable data collection for these data sets, set the value for the **APM\_ATTRIBUTES\_ENABLE\_COLLECTION** environment variable to No.

4. Restart the Microsoft Active Directory agent.

#### **What to do next**

Log in to the Cloud App Management console to view data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

### **Running Microsoft Active Directory agent as a non-administrator user**

You can run the Log File agent as a non-administrator user.

#### **About this task**

You can run the monitoring agent for Active Directory as a non-administrator user; however, Trust Topology attributes and Sysvol Replication attributes might not be available. These attributes are available only to domain users.

To view the Trust Topology attributes, a non-administrator user must have the following registry permissions:

- Grant full access to the HKEY\_LOCAL\_MACHINE\SOFTWARE\Candle directory.
- Grant read access to the HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\WindowsNT\CurrentVersion\Perflib directory.

To view the Sysvol Replication attributes, a non-administrator user must have full access to the Sysvol folder on all domain controllers in a domain.

**Important:** When Microsoft Active Directory agent is running as a non-administrator user, some services from the Services attribute group show values for Current State and Start Type attributes as Unknown on the APM User Interface.

The following table contains the attribute groups for the Active Directory agent that display data for domain users and performance monitoring users.

Table 42. Attribute groups for domain users and performance monitoring users

User right	Attribute group
Domain users	<ul style="list-style-type: none"> <li>• RID Pool Information</li> <li>• Services</li> <li>• Event Logs</li> <li>• DNS</li> <li>• DNS ADIntegrated Details</li> <li>• DNS ADIntegrated</li> <li>• DHCP</li> <li>• Trust</li> <li>• Group Policy Objects</li> <li>• Lost and Found Objects</li> <li>• Exchange Directory Service</li> <li>• Replication Conflict Objects</li> <li>• LDAP Attribute</li> <li>• Root Directory Server</li> <li>• Containers</li> <li>• Replication Partner</li> <li>• Domain Controller Availability</li> <li>• Replication Partner Latency</li> <li>• Forest Topology</li> </ul>
Domain users and performance monitoring users	<p>All attribute groups that are mentioned for the domain users and the following extra attribute groups:</p> <ul style="list-style-type: none"> <li>• Address Book</li> <li>• Replication</li> <li>• Directory Services</li> <li>• Knowledge Consistency Checker</li> <li>• Kerberos Key Distribution Center</li> <li>• Lightweight Directory Access Protocol</li> <li>• Local Security Authority</li> <li>• Name Service Provider</li> <li>• Security Accounts Manager</li> <li>• File Replication Service</li> <li>• Distributed File System Replication</li> <li>• DFS Replication Connections</li> <li>• DFS Replicated Folders</li> <li>• DFS Service Volume</li> <li>• Domain Controller Performance</li> <li>• Remote Access Server</li> <li>• Direct-Access Server</li> <li>• Netlogon Attributes</li> </ul>

**Note:** Additionally, the following attribute groups display data for users who are members of the *Administrators* group:

- Active Directory Database Information
- Moved or Deleted Organizational Unit
- Password Setting Objects

For information, refer [“Configuring Microsoft Active Directory monitoring” on page 340](#)

### Procedure

1. Click **Start > Programs > Administrative Tools > Active Directory Users and Computers**.
2. Expand the domain in which you want to create the user by clicking the plus sign (+) next to the name of a domain.
3. Right-click **Users**, and then click **New > User**.
4. Create a new user by using the **New Object - User** wizard. By default, a new user is a member of the **Domain Users** group.
5. Right-click the new user that is created in the *Domain Users* group, and click **Properties**. The **Username Properties** window opens, where *username* is the name of the new user. Complete the following steps in the **Username Properties** window:
  - a) Click **Member of** tab. In the **Member of** area, add the **Performance Monitor Users** group.
  - b) Click **Apply**, and then click **OK**.
6. Go to the `Candle_Home` directory. The default path is `C:\IBM\APM`.
7. Right-click the APM folder and click **Properties**. The **APM Properties** window opens. Complete the following steps in the **APM Properties** window:
  - a) On the **Security** tab, click **Edit**.
  - b) Click **Add** to add the new user and grant full access to this user.
  - c) Click **Apply**, and then click **OK**.
8. Click **Start > Run**, and then type `services.msc`. The **Services** window opens. Complete the following steps in the **Services** window:
  - a) Right-click the **Monitoring Agent** for Active Directory service, and click **Properties**.
  - b) In the **Active Directory Properties** window, on the **Log On** tab, click **This Account**. Enter the user credentials.
  - c) Click **Apply**, and then click **OK**.
9. Restart the agent service.

## Configuring Microsoft Cluster Server monitoring

You must configure the Monitoring Agent for Microsoft Cluster Server so that the agent can collect the cluster server data. Use the silent response file to configure the agent.

### Before you begin

Ensure that you complete the following tasks:

- Create an empty resource group for the agent.
- Create a generic service cluster resource in the resource group of the agent on Windows Server 2008, 2012, 2016, and 2019 systems.
- You must have administrator privileges to connect to the Microsoft Cluster Server environment or application. You can create a new user and assign administrator privileges by adding the new user to the Administrators group.

**Remember:** To configure the Microsoft Cluster Server agent, you can use a local or a domain user who has administrator privileges.

The product version and the agent version often differ. The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see the [“Change history”](#) on page 52.

### About this task

The Microsoft Cluster Server agent is a single instance agent. You must install and configure the agent manually in the same way on each node in the cluster. To configure the agent, see [“Configuring the agent by using the silent response file”](#) on page 345.

## Creating a generic service cluster resource on Windows Server 2008, 2012, 2016, and 2019 systems

You must add the cluster agent service as a resource for the agent to monitor the cluster server.

### Before you begin

Ensure that the agent is in stopped state on each node in the cluster.

### Procedure

To create a generic service cluster resource, complete the following steps:

1. Open the **Failover Cluster Manager** on any one of the cluster nodes.
2. Complete one of the following steps:
  - For Windows Server 2008:  
In the navigation pane, right-click **Services And Applications**, and then click **More Actions > Create Empty Service or Application**. The new service displays in the services and applications list. Rename the newly created service.
  - For Windows Server 2012:  
In the navigation pane, right-click **Roles**, and then click **More Actions > Create Roles**. The new service is displayed in the roles list.
  - For Windows Server 2016 and 2019:  
In the navigation pane, right-click **Roles**, and then click **Configure Roles**. The new service displays.
3. Right-click the new service and click **Add resource > Generic Service**.
4. In the **New Resource Wizard** window, select **Monitoring Agent for Microsoft Cluster Server** and click **Next**.
5. Click **Next** in the subsequent windows until you see **Finish**.
6. Click **Finish**.  
The agent service is added as a resource.
7. Right-click **Monitoring Agent for Microsoft Cluster Server** resource and click **Bring Resource Online**.

### Results

The agent is started on the preferred node.

## Configuring the agent by using the silent response file

The silent response file contains the Microsoft Cluster Server agent configuration parameters with default values that are defined for some parameters. You can edit the silent response file to configure the agent with different values for the configuration parameters.

### Before you begin

If you want to modify the default configuration parameters, edit the response file.

### About this task

You can configure the agent by using the silent response file.

### Procedure

1. Open the silent response file that is available at following path:  
`install_dir\samples\microsoft_cluster_server_silent_config.txt`
2. Enter the value of **CTIRA\_HOSTNAME** environment variable as the cluster name.
3. On each cluster node, run the following command:

```
install_dir\BIN\microsoft_cluster_server-agent.bat config install_dir\samples  
\microsoft_cluster_server_silent_config.txt
```

### What to do next

Change the user account from the local user to the domain user.

## Changing the user account

After you configure the Microsoft Cluster Server agent, you can change the user account from the local user to the domain user.

### About this task

By default, the agent runs under the local user account. The agent must be run under the domain user so that the agent can monitor all nodes in the cluster from a single node.

### Procedure

To change the user account, complete the following steps:

1. Open the **IBM Performance Management** window.
2. Right-click the agent and click **Change Startup**.
3. Enter the domain login credentials.
4. Open the **Failover Cluster Manager** on one of the nodes, and start the cluster service.

### Results

The agent is started on the node.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Performance Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring Microsoft Exchange Server monitoring

You must configure the Monitoring Agent for Microsoft Exchange Server to monitor the availability and performance of Microsoft Exchange Server.

### Before you begin

- Review the hardware and software prerequisites. For the up-to-date system requirement information, see the Software Product Compatibility Reports (SPCR) for the [Software Product Compatibility Reports \(SPCR\) for the Microsoft Exchange Server agent](#).
- The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history”](#) on page 52.
- Ensure that you complete the following tasks:
  - [“Creating users”](#) on page 347

- [“Assigning administrator rights to the Exchange Server user” on page 349](#)
- [“Making the Exchange Server user a local administrator” on page 350](#)
- [“Configuring the Exchange Server for reachability” on page 351](#)
- [“Configuring the agent to run under the domain user” on page 352](#)

### About this task

You can start the Microsoft Exchange Server agent after the agent is installed. Manual configuration is required to view data for all the agent attributes.

- To configure the agent locally, see [“Configuring the agent locally” on page 353](#).
- To configure the agent by using the silent response file, see [“Configuring the agent by using the silent response file” on page 357](#).

## Creating users

You can create a user for the agent on the Exchange Server manually or by running the *New User* utility. You must create the user on each Exchange Server that you want to monitor.

### Before you begin

Install the Microsoft Exchange Server agent. To create a user, you must be a domain administrator with full administrator rights on the Microsoft Exchange Server.

### About this task

Use one of the following procedures to create users:

- [“Creating users on Exchange Server 2007 and 2010” on page 347](#)
- [“Creating users by running the New User utility” on page 348](#)

### Creating users on Exchange Server 2007 and 2010

You must create a user for the agent on Exchange Server 2007 and 2010 so that the agent can communicate and authenticate with the Exchange Server that you want to monitor.

### Procedure

To create a user, complete the following steps:

1. Click **Start > Programs > Microsoft Exchange Server 2007 > Exchange Management Console**. The **Exchange Management Console** window opens.
2. In the Console tree, click **Mailbox in Recipient Configuration**.
3. In the Action pane, click **New Mailbox**. The New Mailbox wizard opens.
4. On the **Introduction** page, click **User Mailbox**.
5. On the **User Type** page, click **New User**.
6. On the **User Information** page, specify the following information:

#### Organizational unit

By default, the users container in the Active Directory is displayed. Click **Browse** to change the default organizational unit.

#### First name

Type the first name of the user.

#### Initials

Type the initials of the user.

#### Last name

Type the last name of the user.

**Name**

By default, the user's first name, initials, and last name are displayed in this field. You can modify the name.

**User log on name (User Principal Name)**

Type the name that the user must use to log on to the mailbox.

**User log on name (pre-Windows 2000, or earlier)**

Type the user name that is compatible with Microsoft Windows 2000 Server, or earlier.

**Password**

Type the password that the user must use to log on to the mailbox.

**Confirm password**

Retype the password that you entered in the **Password** field.

**User must change password at next logon**

Select this check box if you want the user to reset the password.

7. On the **Mailbox Settings** page, specify the following information:

**Alias**

By default, the value for this field is identical to the value that you specified in the **User logon name (User Principal Name)** field.

**Mailbox database**

Click **Browse** to open the **Select Mailbox Database** window. Select the mailbox database that you want to use and click **OK**.

**Managed folder mailbox policy**

Select this check box to specify a messaging records management (MRM) policy. Click **Browse** to select the MRM mailbox policy that you want to associate with this mailbox.

**Exchange ActiveSync mailbox policy**

Select this check box to specify an Exchange ActiveSync mailbox policy. Click **Browse** to select the Exchange ActiveSync mailbox policy that you want to associate with this mailbox.

8. On the **New Mailbox** page, review the configuration summary. Click **New** to create a mailbox. On the **Completion** page, the Summary section shows whether the mailbox was created.

9. Click **Finish**.

**What to do next**

Assign administrator rights to the Exchange user that you created.

**Creating users by running the New User utility**

You can run the New User utility to create users on Exchange Server 2007 or later. The user that is created by running this utility has all the required permissions to run the agent. This utility is installed when you install the agent.

**Before you begin**

Ensure that the agent is installed. To run the New User utility, you must be a domain administrator with full administrator rights on the Exchange Server.

**About this task**

When you run this utility, the user is created in the Users group of the Active Directory, and has the following permissions:

- On Exchange Server 2007:
  - Local administrator
  - Remote desktop user
  - Exchange recipient administrator
- On Exchange Server 2010, or later:



- Local administrator
- Remote desktop user
- Exchange Servers or Public Folder Management.

## Procedure

To run the New User utility, complete the following steps:

1. Double-click the `kexnewuser1.exe` file that is available at the following location:  
`install_dir\TMAITM6_x64` Where *install\_dir* is the path where the agent is installed.
2. In the **New User** window, complete the following steps:
  - a) Enter the **first name** and the **last name** of the user.  
**Restriction:** The length of the first and the last name must not exceed 28 characters.
  - b) In the **User Logon Name** field, enter the name that the user must type whenever the user logs in.  
**Restriction:** The length of the user logon name must not exceed 256 characters.
  - c) In the **Password** field, enter your password.
  - d) In the **Confirm Password** field, enter the password again.
  - e) Select **User Must Change Password at Next Logon** if you want to reset the specified password when the user logs on next time.
  - f) Click **Next**.

The configuration values that you specify are validated, and error messages are displayed for incorrect values.

3. From the list of mailbox databases, select the required mailbox database, and click **Next**.  
A summary of configuration values is displayed.
4. Click **Finish**.

## Results

The settings are saved, and the user is created.

## Assigning administrator rights to the Exchange Server user

The user that you created for the Microsoft Exchange Server agent must have administrator rights to access the Microsoft Exchange Server components.

### Before you begin

Create an Exchange Server user who has the mailbox on the Exchange Server that is being monitored.

### About this task

Use one of the following procedures to assign administrator rights to the user:

- [“Assigning administrator rights on Exchange Server 2007” on page 349](#)
- [“Assigning administrator rights on Exchange Server 2010” on page 350](#)

### Assigning administrator rights on Exchange Server 2007

You must assign Exchange Recipient Administrator rights to the user on Exchange Server 2007.

## Procedure

1. Click **Start > Programs > Microsoft Exchange Server 2007 > Exchange Management Console**. The **Exchange Management Console** window opens.
2. In the Console tree, click **Organization Configuration**.
3. In the Action pane, click **Add Exchange Administrator**.

4. On the **Add Exchange Administrator** page, click **Browse**. Select the new user that you created, and then select **Exchange Recipient Administrator** role.
5. Click **Add**.
6. On the **Completion** page, click **Finish**.

### **Assigning administrator rights on Exchange Server 2010**

You must assign Exchange Servers or Public Folder Management rights to the user on Exchange Server 2010.

#### **Procedure**

1. Log on to Exchange server with Administrator privileges.
2. Click **Start > Administrative Tools > Server Manager**.
3. Expand **Tools**.
4. Click **Active Directory Users and Computers**.
5. Expand **Domain** and click **Microsoft Exchange Security Groups**.
6. Right-click **Exchange Servers or Public Folder Management** and then click **Properties**.
7. In **Exchange Servers Properties or Public Folder Management Properties** window, go to **Members** and click **Add**.
8. From the list of users, select the user that you want to add to the group, and click **OK**.
9. Click **OK**.

### **Making the Exchange Server user a local administrator**

To access the Exchange Server data, the user that you created for the Microsoft Exchange Server agent must be a local administrator of the computer where the Exchange Server is installed.

#### **Before you begin**

Create an Exchange Server user.

#### **About this task**

Use one of the following procedures to make the user a local administrator:

- [“Making the user a local administrator on Windows 2003 computer” on page 350](#)
- [“Making the user a local administrator on Windows 2008 computer” on page 351](#)
- [“Making the user a local administrator on Windows 2012 computer” on page 351](#)
- [“Making the user a local administrator on Windows 2016 computer” on page 351](#)

#### **Making the user a local administrator on Windows 2003 computer**

You must make the user that you created for the Exchange Server a local administrator of the computer that runs on the Windows 2003 operating system, and where the Exchange Server is installed.

#### **Procedure**

1. Right-click **My Computer** on the computer desktop and click **Manage**.
2. Expand **Local Users and Groups**.
3. Click **Groups**.
4. Double-click **Administrators** to display the **Administrators Properties** window.
5. Click **Add**.
6. Select **Entire Directory** from the **Look in** list.
7. Select the name of the user that you created and click **Add**.
8. Click **OK**.
9. Click **OK**.

### **Making the user a local administrator on Windows 2008 computer**

You must make the user that you created for the Exchange Server a local administrator of the computer that runs on the Windows Server 2008 operating system, and where the Exchange Server is installed.

#### **Procedure**

1. Click **Start > Administrative Tools > Server Manager**.
2. In the navigation pane, expand **Configuration**.
3. Double-click **Local Users and Groups**.
4. Click **Groups**.
5. Right-click the group to which you want to add the user account, and then click **Add to Group**.
6. Click **Add** and type the name of the user account.
7. Click **Check Names** and then click **OK**.

### **Making the user a local administrator on Windows 2012 computer**

You must make the user that you created for the Exchange Server a local administrator of the computer that runs on the Windows Server 2012 operating system and where the Exchange Server is installed.

#### **Procedure**

1. Click **Start > Server Manager**.
2. On the **Server Manager dashboard** page, click **Tools > Computer Management**.
3. In the navigation pane of the **Computer Management** page, expand **Local Users and Groups**, and then click **Users**.
4. From the users list, right-click the user to which you want to assign administrator rights, and click **Properties**.
5. Click the **Member Of** tab, and click **Add**.
6. On the **Select Group** page, type Administrators, and then click **OK**.
7. Click **Apply** and **OK**.

### **Making the user a local administrator on Windows 2016 computer**

You must make the user that you created for the Exchange Server a local administrator of the computer that runs on the Windows Server 2016 operating system and where the Exchange Server is installed.

#### **Procedure**

1. Click **Start > Server Manager**.
2. On the **Server Manager dashboard** page, click **Tools > Computer Management**.
3. In the navigation pane of the **Computer Management** page, expand **Local Users and Groups**, and then click **Users**.
4. From the users list, right-click the user to which you want to assign administrator rights, and click **Properties**.
5. Click the **Member Of** tab, and click **Add**.
6. On the **Select Group** page, type Administrators, and then click **OK**.
7. Click **Apply** and **OK**.

## **Configuring the Exchange Server for reachability**

To verify reachability, the Microsoft Exchange Server agent sends an email message to the server, and measures the amount of time to receive an automated response. Before you start the agent, you must configure the Exchange Server to automatically respond to email messages.

### **Before you begin**

Before you configure the Exchange Server, ensure that the following tasks are completed:

- A mailbox is created for the user on the Exchange Server that you want to monitor.
- The user that you created for the agent is a domain user.
- The servers in your Microsoft Exchange organization are configured for mail flow between servers.

### Procedure

To verify reachability of Exchange Server, complete the following steps for each Exchange Server:

1. Log in to Microsoft Outlook by specifying credentials of the user that you created.
2. Click **Next** on the **Startup** window.
3. Select **Yes** and click **Next**.
4. In the **Microsoft Exchange Server** field, type the name of the Exchange Server.
5. In the **Mailbox** field, type the name of the user that you created.
6. Click **Finish**.
7. Click **OK**.
8. Click **Tools > Rules and Alerts > New Rule**.
9. Select **Start from a blank rule**.
10. Select **Check messages when they arrive** and click **Next**.
11. Select the following options:
  - **Where my name is in the To: box**
  - **With specific words in the subject or body**
12. Under **Step 2** in the window, click **Specific words**.
13. In the **Specify words or phrases to search for in the subject or body** field, type AVAILABILITY CHECK.
14. Click **Add**.
15. Click **OK** and then click **Next**.
16. Select **Have the server reply using a specific message** and click **a specific message**.
17. In the email message editor, type the following text in the subject field of the message:  
CHECK RECEIVED: MAILBOX AVAILABLE .
18. Close the email message editor and click **Yes** to save these changes.
19. Click **Next**.
20. When you are asked about exceptions, do not specify any restrictions.
21. Click **Next**.
22. Click **Finish** and then click **OK**.

### What to do next

Configure the Microsoft Exchange Server agent.

## Configuring the agent to run under the domain user

By default, the Microsoft Exchange Server agent is configured to run under the local user. The agent must be run under the domain user that you created.

### Before you begin

Ensure that:

- The user that you created is a domain user with local administrator rights.
- The user has administrator rights to access the Microsoft Exchange Server components.

### About this task

When the agent is run as the domain user, the agent can monitor all the components of the Exchange Server.

## Procedure

To change the user under which the agent runs, complete the following steps:

1. Run the following command to verify which user ID is being used for starting the agent.  
**install\_dir\InstallITM\KinCinfo.exe -r**
2. If the agent was started with a user ID that does not belong to the Administrator group, stop the agent.
3. Open the **Manage Monitoring Services** window.
4. Right-click the agent instance, and click **Change Startup**.
5. Specify the fully qualified user ID as <Domain\Userid>, and then specify the password.
6. Start the monitoring agent.

## Configuring the agent locally

You can configure the agent locally by using the IBM Cloud Application Performance Management window.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Monitoring Agent for Microsoft Exchange Server**, and then click **Configure agent**.



**Attention:** Click **Reconfigure** if **Configure agent** is disabled.

3. In the **Monitoring Agent for Microsoft Exchange Server: Agent Advanced Configuration** window, click **OK**.
4. In the **Agent Configuration** window, complete the following steps:
  - a) Click the **Exchange Server Properties** tab, and specify values for the configuration parameters. When you click **OK**, the specified values are validated.
  - b) Click the **Exchange Services Monitoring** tab, and specify values for the configuration parameters. When you click **OK**, the specified values are validated.
  - c) Click the **Advanced Configuration Properties** tab, and specify values for the configuration parameters. When you click **OK**, the specified values are validated.

For information about the configuration parameters in each tab of the **Agent Configuration** window, see the following topics:

- [“Configuration parameters for the Exchange Server properties” on page 353](#)
- [“Configuration parameters for Exchange services” on page 355](#)
- [“Configuration parameters for reachability” on page 355](#)

For information about the validation of configuration values, see [“Validation of configuration values” on page 356](#).

5. Recycle the agent.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

### Configuration parameters for the Exchange Server properties

In the **Exchange Server Properties** tab of the **Agent Configuration** window, you can configure the Exchange Server properties, such as server name, domain name, and user name.

The following table contains detailed descriptions of the configuration settings in the **Exchange Server Properties** tab.

Table 43. Names and descriptions of configuration settings in the Exchange Server Properties tab

Parameter name	Description	Mandatory field	Examples
Exchange Server Name	<p>The name of the Exchange Server. During installation of the Exchange Server, the default Exchange Server name is the Windows Server host name. If you change the default Exchange Server name, you must use the changed name when you configure the Exchange Server agent.</p> <p><b>Remember:</b> In clustered and distributed environments, specify the Mailbox Server name for Exchange Server 2007.</p>	<p>Yes</p> <p><b>Important:</b> Do not specify a value if the agent is installed on a server that has a single copy cluster with more than two nodes.</p>	<p>If the Exchange Server name is popcorn, enter popcorn in the <b>Exchange Server Name</b> field.</p>
Exchange Domain Name	<p>The name of the domain where the Exchange Server is installed.</p>	Yes	<p>If the Exchange Server is in the LAB.XYZ.com domain, enter the name that precedes the first dot, for example, LAB.</p>
Exchange User Name	<p>The name of the user who is configured to access the Exchange Server.</p> <p><b>Remember:</b> The user must have a mailbox on the same Exchange Server.</p>	Yes	
Exchange User Password	<p>The password of the user who is configured to access the Exchange Server.</p>	Yes	
Confirm Password	<p>The same password that you specified for the Exchange Server user.</p>	Yes	
Exchange MAPI Profile Name	<p>MAPI profiles are the primary configuration settings that are required for accessing the Exchange Server. This field is disabled if you are using a 64-bit Microsoft Exchange Server agent to monitor Exchange Server 2007, or later.</p>	No	
Configuration in cluster	<p>Select this check box if you want to configure the Microsoft Exchange Server agent in a cluster environment.</p>	Not applicable	
Cluster Server Name	<p>The name of the Cluster Server.</p> <p>This field is enabled when you select the <b>Configuration in cluster</b> check box.</p>	Yes, if the field is enabled.	SCCcluster
Exchange Subsystem ID	<p>The name of the Cluster Server node.</p> <p>This field is enabled when you select the <b>Configuration in cluster</b> check box.</p>	Yes, if the field is enabled.	node1

Table 43. Names and descriptions of configuration settings in the Exchange Server Properties tab (continued)

Parameter name	Description	Mandatory field	Examples
Exchange Agent Historical Data Directory	The location on the disk where the historical data is stored.  This field is enabled when you select the <b>Configuration in cluster</b> check box.	Yes, if the field is enabled.	c:\history

#### Configuration parameters for Exchange services

In the **Exchange Services Monitoring** tab of the **Agent Configuration** window, you can select the Exchange services to know the Exchange Server status.

The following table contains detailed descriptions of the configuration settings in the **Exchange Services Monitoring** tab.

Table 44. Names and descriptions of configuration settings in the Exchange Services Monitoring tab

Parameter name	Description	Mandatory field
Exchange Services	Select the Exchange services from the available list of services, and click the arrow to move the selected services to the <b>Services Configured for Server Status</b> list so that the Microsoft Exchange Server agent can monitor them.  <b>Remember:</b> The list of available services changes according to the Exchange Server version and the roles that are installed.	Not applicable
Services Configured for Server Status	The services that are already available in this list determine the status of the Exchange Server. These services are mandatory and cannot be moved from the <b>Services Configured for Server Status</b> list to the <b>Exchange Services</b> list. You can add more services to the <b>Services Configured for Server Status</b> list by moving the services from the <b>Exchange Services</b> list. You can move these additional services back to the <b>Exchange Services</b> list.	Not applicable

#### Configuration parameters for reachability

In the **Advanced Configuration Properties** tab of the **Agent Configuration** window, you can configure the parameters that are related to reachability, such as target email address and reachability interval.

The following table contains detailed descriptions of the configuration settings in the **Advanced Configuration Properties** tab.

Table 45. Names and descriptions of configuration settings in the Advanced Configuration Properties tab

Parameter name	Description	Mandatory field
Enable Mailbox Reachability Monitoring	Select this check box if you want the agent to capture the reachability metrics data.	Not applicable
Target Email Address	An email address to verify reachability. Separate multiple email addresses with a semicolon (;).  <b>Restriction:</b> The total number of characters in this field must not exceed 1023.	Yes, if this field is enabled.
Email Transmission Interval (seconds)	The waiting time (in seconds) of the Exchange Server agent between sending emails.	Yes, if this field is enabled.

Table 45. Names and descriptions of configuration settings in the Advanced Configuration Properties tab (continued)

Parameter name	Description	Mandatory field
Email Transmission Timeout (seconds)	The interval (in seconds) for which the agent waits for a response to the email that was sent to test whether the Mailbox Server is reachable.	No
Enable Mailbox Detail Monitoring	Select this check box to collect data for the mailbox detail metrics.	Not applicable
Mailbox Detail Collection Start time	The time (in hh:mm:ss format) when mailbox detail metrics are collected.	No
Mailbox Detail Collection Interval (seconds)	The interval (in seconds) between collections of mailbox detail metrics.	No
Event Logs Collection Time (minutes)	The duration (in minutes) for which the agent collects event records.	No
Maximum Number of Events	The maximum count up to which event records are collected. The collection of event records stops when the number of collected event records exceeds the maximum count.	No
Collection Interval (seconds)	The interval (in seconds) between the agent cycles.	No
Exchange Topology Interval (seconds)	The interval (in seconds) between collections of topology detail information.	No
Message Tracking Collection Interval (hours)	<p>The interval (in hours) for which the message tracking logs are collected.</p> <p><b>Restriction:</b> The interval value must be in the range 1 - 12. If you specify the interval value that is greater than 12, the value is saved as 12. If you enter an invalid value that contains alphabets or special characters, the value is saved as 0, which indicates that the message tracking collection is disabled.</p> <p>This field is disabled if any of the following conditions is true:</p> <ul style="list-style-type: none"> <li>The Mailbox Server role or the Hub Transport role is not installed on the Exchange Server.</li> <li>The message tracking feature is disabled on the Exchange Server.</li> </ul>	No

### Validation of configuration values

The values that you specify while configuring the agent are validated. The validation ensures that the values are specified for all mandatory parameters and certain conditions are met, such as local administrator rights for the user.

The following table shows the validation tests that are performed on the specified configuration values.

Table 46. Validation tests	
Validation test	Verifies whether
Exchange Server Name	The Mailbox Server name of the user matches the specified Exchange Server name.



Table 46. Validation tests (continued)

Validation test	Verifies whether
Exchange Server Rights	The user has the required Exchange Server rights. On Exchange Server 2007, the user must have recipient administrator rights, and on Exchange Server 2010, or later, the user must have recipient management rights.
Local Admin	The user has local administrator rights.
Agent Service Logon	The agent service is configured to run with the specified user account.

If one or more validation tests fail, an error message is generated. You must specify values for all mandatory parameters. Otherwise, you cannot save the configured values.

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters with default values defined for some parameters. You can edit the silent response file to configure the agent with different values for the configuration parameters.

### About this task

After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

1. Open the `msex_silent_config.txt` file that is located at `install_dir\samples`, and specify values for all mandatory parameters.

You can also modify the default values of other parameters.

2. Run the following command:

```
install_dir\BIN\msexch-agent.bat config install_dir\samples  
\msex_silent_config.txt
```

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring local environment variables for the agent

You can configure the local environment variables for the Microsoft Exchange Server agent to enable or disable event throttling for duplicate events.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Cloud Application Performance Management**.
2. In the **IBM Performance Management** window, from the **Actions** menu, click **Advanced > Edit ENV File**.
3. In the KEXENV file, change the values of the following environment variables:

#### **EX\_EVENT\_THROTTLE\_ENABLE**

This variable enables you to throttle duplicate events. The default value is `False`. To enable event throttling to prevent triggering of situations for duplicate events, set the value of this variable to `True`.

#### **EX\_EVENT\_THROTTLE\_DURATION**

This variable provides the duration (in minutes) for throttling of events. The default value is 0 minutes.

## Configuring Microsoft Hyper-V monitoring

When you install the Monitoring Agent for Microsoft Hyper-V Server, the agent is automatically configured and started with the default configuration settings. Use the silent response file to modify the default configuration settings.

### Before you begin

- Review hardware and software prerequisites. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\)](#)
- If you want to modify the default configuration parameters, edit the response file.
- To view the virtual machine data in the Virtual Machine page, ensure that you install the integration component and the OS agent on each virtual machine. For virtual machines that run on the Linux system, ensure that you complete the following tasks:
  - Upgrade the Linux system.
  - Install the updated `hyperivkvpd` or `hyperiv-daemons` rpm package on the virtual machine.

### About this task

You can configure the agent when the agent is running or stopped. The agent remains in the same state after configuration. For example, if the agent is running, it remains in the running state after configuration.

### Procedure

To configure the agent, complete the following steps:

1. Open the `microsoft_hyper-v_server_silent_config.txt` file that is at `install_dir\samples`, and specify values for all mandatory parameters.

You can also modify the default values of other parameters.

2. Open the command prompt, and enter the following command:

```
install_dir\BIN\microsoft_hyper-v_server-agent.bat config install_dir\samples\microsoft_hyper-v_server_silent_config.txt
```

The response file contains the following parameters:

- KHV\_DIRECTOR\_PORT
- KHV\_DIRECTOR\_SERVER

**Remember:** The agent configuration is organized into the following groups:

#### **IBM Systems Director configuration (IBM\_DIRECTOR\_CONFIGURATION)**

The configuration elements that are defined in this group are always present in the agent's configuration. This group defines information that applies to the entire agent.

#### **IBM Systems Director Server Port Number (KHV\_DIRECTOR\_PORT)**

The port number for the IBM Systems Director Server. The default value is none.

#### **IBM Systems Director Server Host Name (KHV\_DIRECTOR\_SERVER)**

The host name or IP address of the IBM Systems Director Server that is managing the environment. The default value is none.

3. Start the agent if it is in the stopped state.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Providing Local Security Policy for running Monitoring Agent for Microsoft Hyper-V Server on Windows by a Non-Administrator user

Local security policies are available to run the Monitoring Agent for Microsoft Hyper-V Server on Windows by a non-administrator user.

### About this task

A combination of following two local security policies works to run the Monitoring Agent for Microsoft Hyper-V Server on Windows by a non-administrator user. For the Monitoring Agent for Microsoft Hyper-V Server to start or stop, configure, and verify data, use these two policies.

- Debug Programs
- Log on as Service

Also, following attribute groups need administrator rights to get data on the Cloud App Management console:

- Availability
- Migration
- VM Mig WO Cluster
- VM Storage Migration

Follow the procedure that is given to avail the Local Security permissions for a non-administrator user.

### Procedure

1. Install the Microsoft Hyper-V Server agent as a local administrator.
2. Add the non-administrator user under the `install_dir` directory and provide the following permissions to it:
  - a) Provide full access to the `HKEY_LOCAL_MACHINE\SOFTWARE\IBMMonitoring` registry.
  - b) Provide read access to the non-administrator user in the `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib` registry.
  - c) Provide full access to the non-administrator user in the `install_dir` directory.
3. Go to the **Start** menu and run the `secpol.msc` command to open the Local Security policies.
4. To add a non-administrator user in the policies, refer [“Granting Local Security Policy permissions” on page 359](#).
5. To add a non-administrator user in the Hyper-V Administrator Users group, refer [“Adding a non-administrator user in the Hyper-V administrator users group” on page 361](#).
6. To add a non-administrator user in the Performance Monitor Users group, refer [“Adding a non-administrator user in the Performance Monitor users group” on page 361](#).
7. To modify the DCOM security permission for a non-administrator user, refer [“Modifying DCOM permissions” on page 360](#).
8. Restart the Monitoring Agent for Microsoft Hyper-V Server and verify data on the Cloud App Management console.

## Granting Local Security Policy permissions

To start or stop, configure, and verify data for the Microsoft Hyper-V Server agent, you need to grant permissions to these two local security policies: Debug Programs and Log on as Service.

### Granting Debug Programs permission

#### About this task

To grant the Debug Programs permission, complete the following procedure on Microsoft Hyper-V Server agent.

## Procedure

1. Click **Start > Control Panel > Administrative Tools > Local Security Policy**. The **Local Security Settings** window opens.
2. Click **Local Policies** to expand the list.
3. Click **User Rights Assignment**. The list of user rights opens.
4. Double-click the **Debug Programs** policy. The **Debug Programs Properties** window opens.
5. Click **Add User or Group**. The **Select Users or Groups** window opens.
6. In the **Enter the object names to select** field, enter the user account name to whom you want to assign permissions, and then click **OK**.
7. Click **Apply**, and then click **OK**.

## Granting Log on as Service permission

### About this task

To grant the Log-on as Service permission, complete the following procedure on Microsoft Hyper-V Server agent.

## Procedure

1. Click **Start > Control Panel > Administrative Tools > Local Security Policy**. The **Local Security Settings** window opens.
2. Click **Local Policies** to expand the list.
3. Click **User Rights Assignment**. The list of user rights opens.
4. Double-click the **Log-on as service** policy. The **Log-on as service Properties** window opens.
5. Click **Add User or Group**. The **Select Users or Groups** window opens.
6. In the **Enter the object names to select** field, enter the user account name to whom you want to assign permissions, and then click **OK**.
7. Click **Apply**, and then click **OK**.

## Modifying DCOM permissions

You need to modify DCOM permissions to run the Microsoft Hyper-V Server agent with the non-administrator user access.

### About this task

To modify DCOM permissions, verify that the user has appropriate permissions to start the DCOM server. To modify permissions, complete the following procedure.

## Procedure

1. Using the **Regedit** command, go to the HKCR\CLSID\clsid registry value.  
**Note:** The CLSID value is displayed in the event viewer with the event ID 10016 when you configure the agent with a non-administrator user.
2. In the Registry Editor pane, double-click **Default**.
3. In the **Edit string** dialog box, copy the value data string.
4. Click **Start > Control Panel > Administration Tools > Component Services**.
5. In the **Component Services** window, expand **Component Services > Computers > My Computer**, and double-click **DCOM**.
6. In the DCOM Config pane, locate the copied string (program name), right-click the program name, and then click **Properties**.
7. In the **Properties** window, select the **Security** tab.

8. Under the **Launch and Activation Permissions** group box, select **Customize**, and then click **Edit**. The **Launch and Activation Permissions** window opens.
9. Click **Add**, enter a non-administrator user to the permission list, and click **OK**.
10. Select the **Allow** check box for Local Launch and Local Activation, and then click **OK**.

## Adding a non-administrator user in the Hyper-V administrator users group

You need to add a non-administrator user in the Hyper-V administrator users group to get data on the Cloud App Management console.

### About this task

To add a non-administrator user in the Hyper-V administrator users group, complete the following procedure.

### Procedure

1. Click **Start > Control Panel > Administration Tools > Computer Management**. The **Computer Management** window opens.
2. In the Computer Management (Local) pane, go to **System Tools > Local Users and Groups > Groups**. The list of groups opens.
3. Double-click the **Hyper-V Administrators** group. The **Hyper-V Administrators Properties** window opens.
4. Click **Add**. The **Select Users or Groups** window opens.
5. In the **Enter the object names to select** field, enter the user account name to whom you want to assign permissions, and then click **OK**.
6. Click **Apply**, and then click **OK**.

## Adding a non-administrator user in the Performance Monitor users group

You need to add a non-administrator user in the Performance Monitor users group to get data on the Cloud App Management console.

### About this task

To add a non-administrator user in the Performance Monitor users group, complete the following procedure.

### Procedure

1. Click **Start > Control Panel > Administration Tools > Computer Management**. The **Computer Management** window opens.
2. In the Computer Management (Local) pane, go to **System Tools > Local Users and Groups > Groups**. The list of groups opens.
3. Double-click the **Performance Monitor Users** group. The **Performance Monitor Users Properties** window opens.
4. Click **Add**. The **Select Users or Groups** window opens.
5. In the **Enter the object names to select** field, enter the user account name to whom you want to assign permissions, and then click **OK**.
6. Click **Apply**, and then click **OK**.

## Configuring Microsoft IIS monitoring

When you install the Monitoring Agent for Microsoft Internet Information Services, the agent is automatically configured and starts with the default configuration settings.

### Before you begin

- The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history” on page 52](#).
- Review hardware and software prerequisites. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\)](#) for Microsoft IIS agent.
- Ensure that the user, who connects to the Microsoft Internet Information Server environment or application, has administrator privileges. Use an existing user with administrator privileges, or create a new user. Assign administrator privileges to the new user by adding the new user to the Administrators group.

**Remember:** To configure the Microsoft IIS agent, you can use a local or a domain user provided that the user has administrator privileges.

### About this task

You can configure the agent when the agent is running or stopped. The agent remains in the same state after configuration. For example, if the agent is running, it remains in the running state after configuration.

The product version and the agent version often differ. The directions here are for the most current release of this agent.

To configure the agent, you can either use the **IBM Performance Management** window or the silent response file.

### What to do next

After you configure the agent, you can change the user account from the local user to the domain user. For steps to change the user account, see [“Changing the user account” on page 364](#).

## Configuring the agent on Windows systems

You can configure the Microsoft IIS agent on Windows operating systems by using the **IBM Performance Management** window. After you update the configuration values, you must start the agent to save the updated values.

### About this task

You can configure the agent when the agent is running or stopped. The agent remains in the same state after configuration. For example, if the agent is running, it remains in the running state after configuration.

The Microsoft IIS agent provides default values for some parameters. You can specify different values for these parameters.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Microsoft IIS agent**, and click **Reconfigure**.
3. In the Monitoring Agent for Microsoft Internet Information Services window, complete the following steps:
  - a) On the **HTTP Error Log Configuration** tab, specify a location to save the log file, and click **Next**.

**Note:** By default, this log file is saved at the following location: C : \WINDOWS \system32\LogFiles\HTTPERR. The administrator can change the location of the log file.

b) On the **Site Log Configuration** tab, specify a location to save the log file, and click **OK**.

**Note:** By default, this log file is saved at the following location: C:\inetpub\logs\LogFiles. The administrator can change the location of the log file.

4. In the **Restart of Monitoring Agent for Microsoft IIS** window, click **Yes**.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

If you need help with troubleshooting, see the [IBM Cloud App Management & IBM Cloud Application Performance Management](#) on developerWorks.

## Configuring the agent by using the silent response file

When you install the Microsoft IIS agent, the agent is automatically configured and starts with the default configuration settings. Use the silent response file to modify the default configuration settings.

### Before you begin

If you want to modify the default configuration parameters, edit the response file.

### About this task

You can configure the agent when the agent is running or stopped. The agent remains in the same state after configuration. For example, if the agent is running, it remains in the running state after configuration.

### Procedure

To configure the Microsoft IIS agent, complete the following steps:

1. On the command line, change the path to the directory that contains the `msiis-agent.bat` file.
2. Enter the following command:  
**msiis-agent.bat config** *absolute path to the response file*.

The response file contains the following parameters:

#### **KQ7\_SITE\_LOG\_FILE**

C:\inetpub\logs\LogFiles

#### **KQ7\_HTTP\_ERROR\_LOG\_FILE**

C:\WINDOWS\system32\LogFiles\HTTPERR

**Remember:** The agent configuration is organized into the following groups:

#### **Site Log Configuration (SITE\_LOG)**

This group contains the configuration parameters that are related to the site log file (KQ7\_SITE\_LOG\_FILE). An administrator can specify a location to save the log file. By default, this log file is saved at the location: C:\inetpub\logs\LogFiles

#### **HTTP Error Log Configuration (HTTP\_ERROR\_LOG)**

This group contains the configuration parameters that are related to the HTTP error log file (KQ7\_HTTP\_ERROR\_LOG\_FILE). An administrator can specify a location to save the log file. By default, this log file is saved at the location: C:\WINDOWS\system32\LogFiles\HTTPERR

3. If the agent is in the stopped state, start the agent.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

If you need help with troubleshooting, see the [IBM Cloud App Management & IBM Cloud Application Performance Management](#) on developerWorks.

## Changing the user account

After you configure the Microsoft IIS agent, you can change the user account from the local user to the domain user.

### About this task

By default, the Microsoft IIS agent runs under the local user account.

### Procedure

1. Run the following command to verify which user ID is being used for starting the agent:

```
install_dir\InstallITM\KinCinfo.exe -r
```

2. If the monitoring agent was started with a user ID that does not belong to the Administrator group, stop the agent.
3. Open the **Manage Monitoring Services** window.
4. Right-click the agent instance, and click **Change Startup**.
5. Specify the fully qualified user ID as <Domain\User ID>, and then specify the password.
6. Start the Microsoft IIS agent.

## Configuring Microsoft Office 365 monitoring

You must configure the Microsoft Office 365 agent to monitor the availability and performance of Microsoft Office 365 subscriptions of the organization.

### Before you begin

- The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [“Using agent commands” on page 226](#). For detailed information about the agent version list and what's new for each version, see [“Change history” on page 52](#).
- Review hardware and software prerequisites. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\) for the Microsoft Office 365 agent](#).
- To collect data for Microsoft office 365 users, the following modules must be installed on the Windows operating system where the agent is installed:
  - PowerShell 3.0 or later
  - Microsoft Online Services Sign-In Assistant PowerShell
  - SharePoint Online Management Shell
  - DotNetFrameworkVersion 4.5.2 or later

Microsoft Office 365 agent supports Windows operating system server 2012 (Datacenter, Enterprise & Standard) Editions 64 bit and above.

To configure the Microsoft Office 365 agent agent, you must have administrative privileges along with privileges to enable the remote execution policy of PowerShell. For more information, see [“Enabling remote execution policy for PowerShell” on page 365](#).

- To monitor Skype synthetic transactions, complete the following tasks:
  - Install the Skype 2013 client on the Windows operating system where you want to perform synthetic transactions for Skype.
  - Set the default video device for Skype as a virtual audio-video filter.
- You must have administrator privileges to start the Microsoft Office 365. For a new user, assign administrator privileges by adding it to the Administrators group.



### About this task

You can start the Microsoft Office 365 agent after the agent is installed. However, manual configuration is required to view data for all the agent attributes.

To configure the agent, you can either use the **IBM Performance Management** window or the silent response file.

## Enabling remote execution policy for PowerShell

You must have administrative privileges along with privileges to enable the remote execution policy of PowerShell.

### About this task

Use the following steps to enable the remote executor policy of PowerShell.

**Important:** This is the one time procedure that you need to complete on your system.

### Procedure

1. Open Windows Power Shell command prompt as administrator and run the following command:

```
Install-Module -<module_name> AzureAD
```

Where <module\_name> is the name of the module you want to install on AzureAD.

2. Connect to AzureAD from your Office 365 subscription, run the following command:

```
Connect-AzureAD
```

3. Enter your Office 365 subscription credentials.

## Verifying reachability of configured users

To verify reachability, Microsoft Office 365 agent sends an email message to the configured users, and measures the amount of time that is required to receive an automated response.

### Before you begin

Ensure that the following tasks are completed:

- Configure all the users, which are configured in the Microsoft Office 365 agent mailbox reachability setting, to automatically respond to email messages.
- A mailbox is created for each user on the Exchange Online that you want to monitor.
- The users that you create for the agent must be global Office 365 user.

### Procedure

Complete the following steps for each Exchange Online user account for which you want to verify reachability:

1. Log in to Microsoft Outlook by specifying the user credential that you created.
2. Go to **Tools > Rules and Alerts > New Rule**.
3. In the **Rules wizard** window, under **Start from a blank rule**, click **Apply rule on messages I receive** and click **Next**.
4. Select one of the following options:
  - **From people or public group**
  - **With specific words in the subject**
5. Under **Step 2** in the window, click **people or public group**.
6. In the **Rule address** window, select the user (global administrator) from which the messages are to be received and click **Next**.

7. Under **Step 2** in the window, click **Specific words**.
8. In the **Specify words or phrases to search for in the subject or body** field, enter text as Test Reachability.
9. Click **Add**.
10. Click **OK** and click **Next**.
11. Select **Have the server reply using a specific message** and click **a specific message**.
12. In the email message editor, type the following text in the subject field of the message:  
Test Reachability.
13. In the **To** list, add the global administrator.
14. Close the email message editor and click **Yes** to save these changes.
15. Click **Finish**.
16. Click **Apply** and click **OK**.

#### What to do next

Configure the Microsoft Office 365 agent on the operating system of your choice.

## Configuring the agent on Windows systems

You must configure the Microsoft Office 365 agent on Windows operating systems by using the Microsoft Office 365 agent window. After you update the configuration values, you must start the agent to save the updated values.

#### About this task

You can configure the agent when the agent is running or stopped state. The agent remains in the same state after configuration.

The Microsoft Office 365 agent provides default values for some parameters. You can specify different values for these parameters.

#### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Monitoring Agent for Microsoft Office 365**, and click **Configure Agent**.
3. In the **Monitoring Agent for Microsoft Office 365** window, complete the following steps:
  - a) On the **Office365 Subscription Details** tab, enter the user name and password of the Office 365 global administrator, and click **Next**.
  - b) On the **Synthetic Transaction** tab, enter the list of email addresses that are delimited by semicolons in the **Reachability Email Addresses** field.
  - c) To enable the data collection of Skype QoS metrics, select the **Skype QoS** check box, and click **Next**.
  - d) On the **Mailbox and OneDrive Usage Monitoring** tab, select the duration for the collection interval in hours from the **Collection Interval** list, and click **Next**.
4. In the **Monitoring Agent for Microsoft Office 365** window, click **Yes**.

#### What to do next

- Configure the Skype synthetic transaction utilities to monitor the Skype QoS synthetic transactions. For more information about monitoring the Skype QoS, see [“Monitoring the Skype's quality of service” on page 368](#).
- Change the user account from the local user to the domain user. For more information, see [“Changing the user account” on page 368](#).

- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent by using the silent response file

When you install the Microsoft Office 365 agent, you configure the agent and start it manually. Use the silent response file to configure the custom settings.

### Before you begin

Edit the response file at <CANDLEHOME>\samples to modify the default configuration settings as follows:

*Table 47.*

Fields	Description
KMO_USER_NAME	The user name of the Office 365 global administrator.
KMO_PASSWORD	The password of the Office 365 global administrator.
KMO_MAIL_ADDRESSES1	A list of email addresses to be targeted for verifying mailbox reachability. The list of email addresses must be delimited using semicolons.
KMO_SKYPE	This parameter is used to enable the collection of the Skype QoS synthetic transactions.
KMO_DATA_COLLECTION_DURATION	The duration, in hours, for which the agent waits before it fetches the mailbox and OneDrive usage data.

### About this task

You can configure the agent when the agent is running or stopped state. The agent remains in the same state after configuration.

### Procedure

To configure the Microsoft Office 365 agent, complete the following steps:

1. On the command line, change the path to the directory that contains the `microsoft_office365-agent.bat` file.
2. Run the following command:

```
microsoft_office365-agent.bat absolute path to the response file
```

3. Optional: If the agent is in the stopped state, start the agent.

### What to do next

- Configure the Skype synthetic transaction utilities to monitor the Skype QoS synthetic transactions. For more information about monitoring the Skype QoS, see [“Monitoring the Skype's quality of service” on page 368](#).
- Change the user account from the local user to the domain user. For more information, see [“Changing the user account” on page 368](#).
- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Changing the user account

After you configure the Microsoft Office 365 agent, change the user account from the local user to the domain user.

### About this task

By default, the Microsoft Office 365 agent runs under the local user account.

### Procedure

1. Run the following command to verify the user ID that is used to start the agent:

```
install_dir\InstallITM\KinCinfo.exe -r
```

2. If the agent is started with user ID that does not belong to the Administrator group, stop the agent.
3. Open the **Manage Monitoring Services** window and right-click the agent instance.
4. Click **Change Startup**.
5. Enter the fully qualified user ID as <Domain\User ID> and password.
6. Start the Microsoft Office 365 agent.

## Monitoring the Skype's quality of service

Configure the Skype synthetic transaction utilities, `kmoskypecaller.exe` and `Kmoskypereceiver.exe` to monitor the Skype's quality of service. You can configure the Skype synthetic transaction utilities on the Windows operating system where the Microsoft Office 365 agent is installed or in a distributed environment where the Skype for Business client is configured.

### Before you begin

To perform synthetic transactions, you must update the Skype caller name and the Skype receiver name in the `<CANDLEHOME>\tmaitm6_x64\kmoskypecallerlist.properties` file according to the following format:

```
skype caller = skype receiver
```

For example, `john@xyz.com = alan@xyz.com`

You can add multiple Skype call receivers for a single Skype caller in the following format:

```
skype caller = list of skype receiver
```

For example, `john@xyz.com = alan@xyz.com;bill@xyz.com;chuk@xyz.com`

**Remember:** If you do not want to perform synthetic transactions but want to monitor the Skype's quality of service for real-time users, do not update the `kmoskypecallerlist.properties` file at `<CANDLEHOME>\TMAITM6_x64` path.

### About this task

When the Microsoft Office 365 agent is configured and started, the following files and folders are created at `<CANDLEHOME>\TMAITM6_x64\`:

- `kmoskypecaller.properties`
- `kmoskypecallerlist.properties`
- `KMOSynthTransSkype.zip`
- `KMOSkypeTransReceiver.zip`

The `kmoskypecaller.properties` file is updated with the server IP and port that is used for communication between the agent and the `kmoskypecaller` utility.

### Procedure

To configure the Skype caller and Skype receivers and initiate synthetic transactions, such as instant messaging, audio and video calls, and application sharing sessions, complete the following steps:

1. Start the Microsoft Office 365 agent.
2. Copy the `KMOSynthTransSkype.zip` file from the agent system to the Windows operating system from where the Skype call is to be initiated.
3. Extract the `KMOSynthTransSkype.zip` file.
4. Copy the `kmoskypecaller.properties` file from the agent system to the extracted `KMOSynthTransSkype` folder on the Windows operating system from where the Skype call is to be initiated.
5. Copy the `KMOSkypeTransReceiver.zip` file from the agent system on all Windows operating system where Skype calls must be received.
6. Extract the `KMOSkypeTransReceiver.zip` file on all Windows operating systems where the Skype calls must be received, and run `KMOSkypeTransReceiver.exe` to start receiving messages.
7. To initiate the synthetic transactions, run the `KMOSynthTransSkype.exe` file that is available in the extracted `KMOSynthTransSkype` folder on the Windows operating system. The Microsoft Office 365 agent starts receiving the Skype monitoring data from the caller client.

### Results

The Microsoft Office 365 agent can now monitor the Skype's quality of service.

## Configuring local environment variables

You can configure local environment variables to change the behavior of the Microsoft Office 365 agent.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. On the **IBM Performance Management** window, from the **Actions** menu, click **Advanced > Edit ENV File**.
3. In the **environment variable** field, enter the values for the environment variables.  
For more information, see [“Local environment variables” on page 369](#).

### Local environment variables

You can change the behavior of the Microsoft Office 365 agent by configuring the local environment variables.

### Variables for defining the data collection method for the agent

To set the method for data collection of the agent, use the following environment variables:

<i>Table 48. Agent data collection</i>	
Environment variables	Description
<b>CDP_DP_INITIAL_COLLECTION_DELAY</b>	Use this variable to set the time interval, in seconds, after which the thread pool begins its data collection
<b>KMO_MAILBOX_REACHABILITY_INTERVAL</b>	Use this variable to set the data collection interval, in minutes, for mailbox reachability attribute group
<b>KMO_SKYPE_REPORT_INTERVAL</b>	Use this variable to set the data collection interval, in hours, for Skype for Business usage statistics feature.
<b>KMO_SERVICE_API_INTERVAL</b>	Use this variable to set the data collection interval, in minutes, for Office 365 service health feature.
<b>KMO_NETWORK_CONNECTION_INTERVAL</b>	Use this variable to set the data collection interval, in minutes, for internet connectivity feature.

Table 48. Agent data collection (continued)

Environment variables	Description
<b>KMO_NETWORK_PERFORMANCE_INTERVAL</b>	Use this variable to set the data collection interval, in minutes, for Office 365 services network performance feature.
<b>KMO_SITE_CONNECTION_INTERVAL</b>	Use this variable to set the data collection interval, in minutes, for Office 365 connectivity feature.
<b>KMO_SPSITE_COLLECTION_INTERVAL</b>	Use this variable to set the data collection interval, in minutes, for SharePoint Sites details feature.
<b>KMO_UASGE_STATS_INTERVAL</b>	Use this variable to set the data collection interval, in hours, for Office 365 Services usage and user statistics feature.
<b>KMO_TENANT_INTERVAL</b>	Use this variable to set the data collection interval, in minutes, for Office 365 tenant details feature.
<b>KMO_ONEDRIVE_CONNECTIVITY_INTERVAL</b>	Use this variable to set the data collection interval, in minutes, for Office 365 OneDrive connectivity feature.
<b>KMO_TENANT_DOMAIN</b>	Use this variable to set the domain name of the tenant.

## Configuring Microsoft SharePoint Server monitoring

When you install the Monitoring Agent for Microsoft SharePoint Server, the agent is automatically configured and started with the default configuration settings. You can use the silent response file to modify the default configuration settings.

### Before you begin

- Review hardware and software prerequisites. For the up-to-date system requirement information, see “[System requirements](#)” on page 75 for Microsoft SharePoint Server agent.
- Ensure that the user, who connects to the Microsoft SharePoint Server environment or application, has administrator privileges. You can use an existing user with administrator privileges, or create a new user. Assign administrator privileges to the new user by adding the new user to the Administrators group.

**Remember:** To configure the Microsoft SharePoint Server agent, you can use a local or a domain user provided the user has administrator privileges.

- Edit the response file and modify the default configuration parameters. The response file contains the following parameters:
  - KQP\_DB\_User  
The user ID of the database.
  - KQP\_DB\_Password  
The password of the database.

The product version and the agent version often differ. The directions here are for the most current release of this agent. For more information about how to check the version of an agent in your environment, see [Agent version](#).

### Procedure

To configure the Microsoft SharePoint Server agent, complete the following steps:

1. Open command prompt and change the path to the directory that contains the `ms_sharepoint_server-agent.bat` file.
2. Enter the following command: `ms_sharepoint_server-agent.bat config absolute path to the response file`.
3. If the agent is in the stopped state, start the agent.

### What to do next

After you configure the agent, you can change the user account from the local user to the domain user. For steps to change the user account, see [“Changing the user account” on page 371](#).

## Changing the user account

After you configure the Microsoft SharePoint Server agent, you can change the user account from the local user to the domain user.

### About this task

With the domain user, the agent can monitor all the components of the Microsoft SharePoint Server agent.

### Procedure

To change the user account, complete the following steps:

1. Run the following command to verify which user ID is being used for starting the agent:  
`install_dir\InstallITM\KinCinfo.exe -r`
2. If the monitoring agent was started with a user ID that does not belong to the Administrator group, stop the agent.
3. Open the **Manage Monitoring Services** window.
4. Right-click the agent instance, and click **Change Startup**.
5. Specify the fully qualified user ID as `<Domain\Userid>`, and then specify the password.
6. Start the monitoring agent.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## Running Monitoring Agent for Microsoft SharePoint Server by a non-admin user

You can use local security policies to run the Monitoring Agent for Microsoft SharePoint Server as a non-administrator user.

### About this task

A combination of following two local security policies works to run the Microsoft SharePoint Server agent by a non-administrator-user:

1. Debug programs
2. Log on as a service

Follow the procedure that is given to avail the Local Security permissions for a non-administrator user.

### Procedure

1. Go to TEMA and change the Microsoft SharePoint Server agent startup with non-administrator user.
2. Add a non-administrator user under the Registry key `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office Server` directory and give read access to it.
3. Add non-administrator user under Registry key `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Shared Tools\Web Server Extensions` and give read access to it.

4. Add non-administrator user manually under Registry key HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Shared Tools\Web Server Extensions\16.0\Secure\ and give read access to it.
5. Add non-administrator user under Registry key HKEY\_LOCAL\_MACHINE\SOFTWARE\IBMMonitoring directory and give full permissions to it.
6. Add non-administrator user under Registry key HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib directory and give read access to it.
7. Add non-administrator user in SharePoint Agent installation folder. For example, C:\IBM\APM and give full permissions to it.
8. Run the **secpol.msc** command in **startmenu** to open the **Local Security Policy**.
9. Add non-administrator user in Local Security Policy. For more information, see [“Local Security Policy permissions”](#) on page 372.
10. Add non-administrator user in the SQL Server Login user group. The user must have sysadmin SQL Server role permissions on the SQL Server.
11. Restart the Microsoft SharePoint Server agent.
12. Check Microsoft SharePoint Server agent status and verify the data on IBM Cloud Application Management portal.
13. The following attribute groups show data for users who are members of the Administrators group:
  - a) Availability

## Local Security Policy permissions

Local security policies are available to run a Microsoft SharePoint Server agent by a non-admin user. These policies help to start or stop, configure, and do data verification of the agent. Following two local security policies work to run the Microsoft SharePoint Server agent by a non-admin-user.

### Granting Log on as a Service permission

You can grant the Log on as a service permission to run the Microsoft SharePoint Server agent as a non-administrator user.

#### About this task

To grant the Log-on as service permission, follow the procedure that is described here.

#### Procedure

1. Click **Start > Administrative Tools > Local Security Policy**. The **Local Security Settings** window opens.
2. In the navigation pane, expand **Local Policy** and click **User Rights Assignment**. The list of user rights opens.
3. Double-click **Log-on as service** policy. The **Log-on as service Properties** window opens.
4. Click **Add User or Group**. The **Select Users or Groups** window opens.
5. In the **Enter the object names to select** field, enter the user account name to whom you want to assign permissions, and then click **OK**.
6. Click **OK**.

### Granting Debug Programs permission

You can grant the Debug Programs permission to run the Microsoft SharePoint Server agent as a non-administrator user..

#### About this task

To grant the Debug Programs permission, follow these steps:



## Procedure

1. Click **Start > Administrative Tools > Local Security Policy**. The **Local Security Settings** window opens.
2. Expand **Local Policy** and click **User Rights Assignment**. The list of user rights opens.
3. Double-click **Debug Programs** policy. The **Debug Programs Properties** window opens.
4. Click **Add User or Group**. The **Select Users or Groups** window opens.
5. In the Enter the object names to select field, enter the user account name to whom you want to assign permissions, and then click **OK**.
6. Click **OK**.

## Configuring Microsoft SQL Server monitoring

You must configure the Monitoring Agent for Microsoft SQL Server so that the agent can collect data from the application that is being monitored.

### Before you begin

The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see ["Change history"](#) on page 52.

Review the hardware and software prerequisites. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\)](#) for the Microsoft SQL Server agent.

You can install and configure the Microsoft SQL Server agent locally by using the command prompt interface. Ensure that the agent is installed on the server that is being monitored.

### About this task

The Microsoft SQL Server agent is a multiple instance agent. You must configure and start each agent instance manually.

- To configure the agent, complete the following tasks:
  - Create a user and grant the required permissions
  - Select the databases for monitoring
  - Configure the local environment variables
- To run the agent in a cluster environment, complete the steps that are described in the "Running the agent in a cluster environment" topic.

## Creating a user and granting permissions

On the Microsoft SQL Server, you must create a user for the agent, and grant permissions to the user for monitoring Microsoft SQL Server. The process of granting permissions is the same for Microsoft SQL Server 2005, or later.

### Before you begin

Install the Microsoft SQL Server agent. To create a user and grant permissions to the user, you must be a database administrator with the **sysdaminn** authorization role.

### About this task

Use the following procedure to determine if an existing SQL Server user has sufficient permissions to monitor Microsoft SQL Server:

- ["Checking the permissions of an existing SQL Server user"](#) on page 374

Use one of the following procedures to create a user:

- ["Creating an SQL Server user ID with Windows authentication"](#) on page 375

- [“Creating an SQL Server user ID with SQL Server authentication” on page 375](#)

Use the following procedure to grant permissions:

- [“Granting minimum permissions for data collection” on page 376](#)
- [“Granting permission to the Perflib registry key for collecting data for few data sets” on page 377](#)

### Checking the permissions of an existing SQL Server user

You can run the utility tool **koqVerifyPermissions.exe** to check if an existing SQL Server user has sufficient permissions related to SQL Server databases.

#### About this task

The utility tool **koqVerifyPermissions.exe** returns the message PASS if the user has **sysadmin** role or the minimum required permissions. The detailed checking result is logged in `koqVerifyPermissions_log`.

The following lists the minimum permissions:

- Permissions for server must include **View server state**, **View any database** and **View any definition**.  
These server level permissions are mandatory.
- For all system databases and the user-defined databases for monitoring, the database role membership must include **public** and **db\_owner**.

The **db\_owner** permission is required to collect data for the following data sets:

- Server details data set
- Database Details data set
- Database Mirroring data set
- Server Summary data set
- Job Summary data set
- For **msdb** database, the database role membership must include **db\_datareader**, **SQLAgentReaderRole** and **SQLAgentUserRole**. These permissions are required for Job Details data set.

#### Procedure

1. Launch the command prompt and change to the following utility directory.
  - For 64-bits agents, `Agent_home\TMAITM6_x64`
  - For 32-bits agents, `Agent_home \TMAITM6`
 where `Agent_home` is the agent installation directory.
2. Run the **koqVerifyPermissions.exe** by providing the parameters:

```
koqVerifyPermissions.exe -S Instance_name -U Username -P Password
```

Where:

- `Instance_name` is the SQL Server instance name.
- `Username` is the user name that is verified by the utility tool.
- `Password` is the password of the user. This parameter is required if `username` is provided.

**Note:** If the `username` and the `password` are not provided, the default user that is logon to the system is used. Example: `NT AUTHORITY\SYSTEM`.

#### Results

The detailed checking result is available in `koqVerifyPermissions_log` at the following directory:

- For 64-bits agents, `Agent_home\TMAITM6_x64\logs`

- For 32-bits agents, *Agent\_home* \TMAITM6\logs

Where *Agent\_home* is the agent installation directory.

### Creating an SQL Server user ID with Windows authentication

Create a new user with the Windows authentication and assign the required roles and permissions to the user.

#### Procedure

To create a user, complete the following steps:

1. In the **SQL Server Management Studio**, open **Object Explorer**.
2. Click **Server\_instance\_name** > **Security** > **Logins**.
3. Right-click **Logins** and select **New Login**.
4. On the **General** page, in the **Login name** field, enter the name of a Windows user.
5. Select **Windows authentication**.
6. Depending on the role and permissions that you want to assign to the user, complete one of the following tasks:
  - On the **Server Roles** page, assign the **sysadmin** role to the new login ID.
  - If you do not want to assign the **sysadmin** role to the user, grant minimum permissions to the user by completing the steps in [“Granting minimum permissions for data collection”](#) on page 376.

**Important:** By default, the **public** role is assigned to the new login ID.

7. Click **OK**.

#### Results

A user is created with the default **public** role and the permissions that you assigned, and is displayed in the **Logins** list.

### Creating an SQL Server user ID with SQL Server authentication

Create a user with the SQL Server authentication and assign the required roles and permissions to the user.

#### Procedure

To create a user, complete the following steps:

1. In the **SQL Server Management Studio**, open **Object Explorer**.
2. Click **Server\_instance\_name** > **Security** > **Logins**.
3. Right-click **Logins** and select **New Login**.
4. On the **General** page, in the **Login name** field, enter the name for a new user.
5. Select **SQL Server authentication**.
6. In the **Password** field, enter a password for the user.
7. In the **Confirm Password** field, retype the password that you entered in the **Password** field.
8. Depending on the role and permissions that you want to assign to the user, complete one of the following tasks:
  - On the **Server Roles** page, assign the **sysadmin** role to the new login ID.
  - If you do not want to assign the **sysadmin** role to the user, grant minimum permissions to the user by completing the steps in [“Granting minimum permissions for data collection”](#) on page 376.

**Important:** By default, the **public** role is assigned to the new login ID.

9. Click **OK**.

## Results

A user is created with the default **public** role and the permissions that you assigned, and is displayed in the **Logins** list.

## Granting minimum permissions for data collection

Apart from the default **public** role, you can assign the **sysadmin** role to a user or grant the minimum permissions to a user so that the agent can collect data for data sets.

## About this task

You can grant the permissions by using the user interface or the utility tool **permissions.cmd**.

## Procedure

- To grant the minimum permissions to the user by using the user interface, complete the following steps:
  - a) Open the **Server Roles** page and verify that the **public** check box is selected.
  - b) Open the **User Mapping** page, and then select the following check boxes for all the system databases and the user-defined databases that you want to monitor:
    - **public**
    - **db\_owner**For the **msdb** database, select the following check boxes:
    - **db\_datareader**
    - **SQLAgentReaderRole**
    - **SQLAgentUserRole**
  - c) Open the **Securables** page, and select the following check boxes for the server instance that you are monitoring:
    - view database
    - view definition
    - view server state
- To grant the minimum permissions to the user by using the utility tool **permissions.cmd**, complete the following steps:
  - a) Start the **Windows Explorer** and browse to the utility tool directory *Agent\_grant\_perm\_dir*:
    - For 64-bits agent, *Agent\_grant\_perm\_dir* is *Agent\_home*\TMAITM6\_x64\scripts\K0Q\GrantPermission.
    - For 32-bits agent, *Agent\_grant\_perm\_dir* is *Agent\_home*\TMAITM6\scripts\K0Q\GrantPermission.
    - The *Agent\_home* is the agent installation directory.



**Attention:** The utility tool **permissions.cmd** grants **db\_owner** on all databases by default. To exclude certain databases, you must add the database names in the *Agent\_grant\_perm\_dir*\exclude\_database.txt file. The database names must be separated by the symbol alias @.

**Tip:** For example, you want to exclude the databases **MyDatabase1** and **MyDatabase2**, add the following entries in the *exclude\_database.txt* file:

```
MyDatabase1@MyDatabase2
```

- b) Double-click **permissions.cmd** to start the utility tool.
- c) Enter the intended parameter values when prompted:

<i>Table 49. Parameters</i>	
<b>Parameters</b>	<b>Description</b>
SQL Server name or SQL Server instance name	Enter the target SQL Server name or the target SQL Server instance name that you want to grant permissions to the user.
The existing SQL Server user's logon name	Enter the user name whose permissions to be altered.
Permissions options: <b>1</b> Grant <b>db_owner</b> permission <b>2</b> Grant <b>db_datareader</b> , <b>SQLAgentReaderRole</b> and <b>SQLAgentUserRole</b> permissions <b>3</b> Grant all required permissions	Enter <b>1</b> or <b>2</b> or <b>3</b> according to your requirement.
The user to grant permissions: <b>1</b> The user who is logon to the system <b>2</b> Another user	Enter <b>1</b> or <b>2</b> . If <b>2</b> is selected, enter the target user name when prompted. <b>Note:</b> The users must have access to grant permissions to other users.

### What to do next

Configure the agent.

### Granting permission to the Perflib registry key for collecting data for few data sets

To collect data for few date sets, you need to grant users read access to the **Perflib** registry key.

### About this task

You need to grant the permission to the Windows user with which agent services are configured. There are many data sets that are affected in absence of **Perflib** permissions. The affected data sets are MS SQL Database Detail, MS SQL Memory Manager, MS SQL Lock Resource Type Summary, MS SQL Job Summary, MS SQL Server Transactions Summary, MS SQL Server Summary and others.

### Procedure

To grant permission to the **Perflib** registry key, complete the following steps:

1. To open Registry Editor, click **Start > Run > Regedit.exe**, and press **Enter**.
2. Go to the HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib registry key.
3. Right-click the **Perflib** key, and click **Permissions**.
4. Click **Add**, enter the windows user name with which the agent is installed and configured, and then click **OK**.
5. Click the user that you added.
6. Allow read access to the user by selecting the check box.
7. Click **Apply**, and then click **OK**.

## Running as a non-administrator user

You can run the monitoring agent for Microsoft SQL Server as a non-administrator user.

### About this task

The Microsoft SQL Server agent can be run as a non-administrator user from Domain Users group.

### Procedure

1. Start Windows application Active Directory Users and Computers and create a domain user.
  - Make sure that the new user is a member of the *Domain Users* group.
  - Make sure that the SQL Server is a member of *Domain Computers*.
2. Add the newly created domain user in the *SQL Server Login* user group. The domain user should have **sysadmin** SQL Server role permission on the SQL Server. For more information, see the [Creating a user and granting permissions](#) topic in the IBM Cloud Application Performance Management Knowledge Center.
3. Log on to the SQL Server as the domain administrator.
4. Grant **Modify** permission to every drive that the Microsoft SQL Server agent accesses. Complete the following procedures to propagate the permission to all sub directories:
  - a) Go to **My Computer**.
  - b) Right-click the **drive**.
  - c) Click the **Security** tab.
  - d) Add the newly created user.
  - e) Give **Modify** permission to the newly created user.
  - f) Click **OK**. This procedure takes a few minutes to apply permission to all sub directories.
5. By using the Windows Registry, grant read access to HKEY\_LOCAL\_MACHINE, and propagate the settings. Complete the following steps to propagate the settings:
  - a) Right-click the HKEY\_LOCAL\_MACHINE directory and select **Permissions**.
  - b) Add the newly created user.
  - c) Select the newly created user.
  - d) Select the **Allow Read** check box.
  - e) Click **OK**. This procedure takes a few minutes to propagate the settings to the entire HKEY\_LOCAL\_MACHINE tree.
6. By using the Windows Registry, grant the agent-specific registry permissions according to the following list.
  - If you installed a 32-bit agent on a 32-bit operating system, grant full access to the KEY\_LOCAL\_MACHINE\SOFTWARE\IBMMonitoring directory, and then propagate the settings.
  - If you installed a 32-bit agent on a 64-bit operating system, grant full access to the HKEY\_LOCAL\_MACHINE\SOFTWARE\Wow6432Node\Candle directory, and then propagate the settings.
  - If you installed a 64-bit agent on a 64-bit operating system, grant full access to the KEY\_LOCAL\_MACHINE\SOFTWARE\IBMMonitoring directory, and then propagate the settings.

Complete the following steps to propagate settings:

- a) Right-click the directory for which you have full access and select **Permissions**.
- b) Add the newly created user.
- c) Select the newly created user.
- d) Select the **Allow Full Control** check box.
- e) Click **OK**. This procedure takes a few minutes to propagate the settings to the entire KEY\_LOCAL\_MACHINE\SOFTWARE\IBMMonitoring tree.

7. Add a new Domain User to the **Performance Monitor Users** group.
8. Verify that Domain Users are members of the *Users* group.
9. Grant the following permissions to the Windows directory to run as a non-administrator user:
  - If a 32-bit agent is installed on a 32-bit operating system, grant read and write access to the `OS_installation_drive:\Windows\system32` directory
  - If a 32-bit agent is installed on a 64-bit operating system, grant read and write access to the `OS_installation_drive:\Windows\SysWOW64` directory

**Note:** Permissions for Windows directory are not necessary for Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012, Windows Server 2012 R2, Windows Server 2016.
10. Grant **Modify** permission to the SQL Server data file and log file:
  - The default path of the SQL Server data file is `SQLServer_root_dir\DATA`, where `SQLServer_root_dir` is the root directory of the SQL Server instance. For example, if the root directory of the SQL Server instance is `C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL`, the data file path is `C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\DATA`.
  - The default path of the SQL Server log file is `SQLServer_root_dir\LOG`, where `SQLServer_root_dir` is the root directory of the SQL Server instance. For example, if the root directory of the SQL Server instance is `C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL`, the log file path is `C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\LOG`.
11. Grant full permissions to the `Candle_Home` directory. The default path is `C:\IBM\ITM`.
12. Apply local security permissions by referring to [“Local Security Policy permissions” on page 379](#).
13. Restart the SQL Server to ensure that local security permissions are applied effectively.
14. Change the logon settings for the SQL Server agent services to the non-administrator user by completing the following steps:
  - a) Click **Start > Administrative Tools > Services**.
  - b) Right-click the **Monitoring Agent For SQL Server *instance\_name***, and click **Properties**. The **SQL Service Properties** window opens.
  - c) Click **Log On** tab.
  - d) Click **This account** and type the user name.
  - e) In the **Password** and **Confirm Password** fields, enter the password, and click **OK**.
  - f) Repeat steps b to e for the **Monitoring Agent For SQL Server Collector *instance\_name***, where *instance\_name* is the Microsoft SQL Server instance name.

### Local Security Policy permissions

Local security policy administers the system and its security policy. It plays an important part in keeping the agent and the system in which the agent is installed secure. This policy works by giving access rights, permissions to users. For, Microsoft SQL Server agent, make sure that the user has following permissions to adhere to local security permission policy.

### [Log on as Service permission](#)

#### About this task

To grant the Log-on as service permission, complete the following steps.

#### Procedure

1. Click **Start > Administrative Tools > Local Security Policy**. The **Local Security Settings** window opens
2. Click **Local Policies** to expand the list.

3. Click **User Rights Assignment**. The list of user rights opens.
4. Double-click **Log-on as service** policy. The **Log-on as service Properties** window opens.
5. Click **Add User or Group**. The **Select Users or Groups** window opens.
6. In the **Enter the object names to select** field, enter the user account name to whom you want to assign the permissions, and click **OK**.
7. Click **OK**.

### *Debug Programs Permission*

#### **About this task**

To grant the debug program permission, complete the following procedure on Microsoft SQL Server agent .

#### **Procedure**

1. Click **Start > Administrative Tools > Local Security Policy**. The **Local Security Settings** window opens.
2. Click **Local Policies** to expand the list.
3. Click **User Rights Assignment**.The list of user rights opens.
4. Double-click **Debug Programs** policy. The **Debug programs Properties** window opens.
5. Click **Add User or Group**. The **Select Users or Groups** window opens.
6. In the **Enter the object names to select** field, enter the user account name to whom you want to assign permissions, and click **OK**
7. Click **OK**.

### *Impersonate a client after authentication*

#### **About this task**

To grant the Impersonate a client after authentication permission, complete the following procedure on Microsoft SQL Server agent .

#### **Procedure**

1. Click **Start > Administrative Tools > Local Security Policy**. The **Local Security Settings** window opens.
2. Click **Local Policies** to expand the list.
3. Click **User Rights Assignment**.The list of user rights opens.
4. Double-click **Impersonate a client after authentication** policy. The **Impersonate a client after authentication Properties** window opens.
5. Click **Add User or Group**. The **Select Users or Groups** window opens.
6. In the **Enter the object names to select** field, enter the user account name to whom you want to assign permissions, and then click **OK**
7. Click **OK**.

## **Selecting the databases for monitoring**

You can select the database that you want to monitor by using the **Configure Database Agents** window.

#### **Procedure**

1. Open the **IBM Performance Management** window.
2. In the **IBM Performance Management** window, click the **Task/SubSystem** column, right-click **Template** and select **Configure Using Defaults**.



3. In the **Configure Database Agents** window, select the database server that you want to monitor from the **Database Servers Available**, and move it to the **Server to Monitor** list.
4. In the **Database Server Properties** window, values for the following fields are automatically populated:
  - Server Name
  - Database Version
  - Home Directory
  - Error Log File

The following fields in the **Database Server Properties** window are optional:

- Windows Authentication
- Support Long Lived Database Connections
- Extended Parms
- Monitor all Databases
- Day(s) Frequency
- Weekly Frequency
- Monthly Frequency
- Collection Start Time
- Table Detail Continuous Collection

For more information about the configuration parameters in the **Database Server Properties** window, see [“Configuration parameters for the Database Server properties”](#) on page 382.

5. If you do not select the **Windows Authentication** field, enter your user ID and password in the **Login** and **Password** fields by using only ASCII characters.
6. In the **Extended Parms** field, enter the name of the data set to disable the data collection, and then click **OK**.

For example:

- Enter koqtbld to disable data collection for Table Detail data set.
  - Enter koqdbd to disable data collection for Database Detail data set.
  - Enter koqtbld, koqdbd to disable data collection for Table Detail and Database Detail data sets.
7. If you do not select the **Monitor All Databases** check box, specify the list of databases for which you want to enable or disable monitoring in the field of **Databases** group area.  
**Remember:** If you select the **Monitor All Databases** check box and specify the databases in **Databases** group area, the setting of **Monitor All Databases** check box takes precedence.
  8. Specify the frequency for the collection of the MS SQL Table Detail data set. The possible values are daily, weekly, or monthly.
  9. Select the **Table Detail Continuous Collection** check box to enable continuous collection of the MS SQL Table Detail data set. If you select the **Table Detail Continuous Collection** check box, enter a value in the **Interval Between Two Continuous Collection (in minutes)** field.
10. In the **Configure Database Agents** window, click **OK**, and then start the agent.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

### Configuration parameters for the Database Server properties

In the **Database Server Properties** window, you can configure the Database Server properties, such as server name, database version and home directory.

The following table contains the descriptions of the configuration settings in the **Database Server Properties** window.

Parameter name	Description	Mandatory field	Examples
Server Name	<p>The name of the Microsoft SQL Server instance that is to be monitored.</p> <p>Use MSSQLSERVER as the instance name for the default instance.</p> <p>The name must be 2 - 32 characters in length to fit the total managed system name.</p>	Yes	<p>If the Microsoft SQL Server instance that is monitored is the default Microsoft SQL Server instance, enter MSSQLSERVER in this field.</p> <p>If the Microsoft SQL Server instance that is monitored is a named instance where the instance name is <code>mysqlserver</code> and the host name is <code>popcorn</code>, enter <code>mysqlserver</code> in this field.</p>
Login	<p>The Microsoft SQL Server user ID used to connect to the Microsoft SQL Server.</p> <p>The user ID is required when <b>Windows Authentication</b> parameter is set to <code>False</code>.</p> <p>Use only ASCII characters for the User ID.</p> <p>When you configure the Microsoft SQL Server agent by specifying a login ID in the <b>Login</b> field, the agent uses this login ID to connect to the Microsoft SQL Server.</p> <p><b>Important:</b> While configuring the agent, if you select the <b>Windows Authentication</b> check box and specify a login ID in the <b>Login</b> field, the agent gives preference to the Windows Authentication.</p>	No	
Password	<p>The password for the Microsoft SQL Server user ID.</p> <p>Password is required only when <b>Windows Authentication</b> parameter is set to <code>False</code>.</p> <p>Use only ASCII characters for the password.</p>	No	

Table 50. Names and descriptions of configuration settings in the **Database Server Properties** window (continued)

Parameter name	Description	Mandatory field	Examples
Database Version	The version of SQL server instance.	Yes	<p>The database versions for SQL server instance are as follows:</p> <ul style="list-style-type: none"> <li>• Microsoft SQL Server 2014 - 12.0.2000.8</li> <li>• Microsoft SQL Server 2012 - 11.0.2100.60</li> <li>• Microsoft SQL Server 2008 R2 - 10.50.1600.1</li> <li>• Microsoft SQL Server 2008 - 10.0.1600.22</li> <li>• Microsoft SQL Server 2005 - 9.0.1399.06</li> </ul>
Home Directory	The SQL server installation directory.	Yes	<p>The default home directory path for the default Microsoft SQL Server 2005 instance is C:\Program Files\Microsoft SQL Server\MSSQL.</p> <p>A named Microsoft SQL Server 2005 instance has a default home directory path in the format C:\Program Files\Microsoft SQL Server\MSSQL \$instance_name, where <i>instance_name</i> is the Microsoft SQL Server instance name.</p>

Table 50. Names and descriptions of configuration settings in the **Database Server Properties** window (continued)

Parameter name	Description	Mandatory field	Examples
Error Log File	The fully qualified location and name of the SQL Server error log.	Yes	<p>The default error log path for the default Microsoft SQL Server 2005 instance is            C:\Program Files\Microsoft SQL Server\MSSQL\LOG\ERRORLOG.</p> <p>A named Microsoft SQL Server 2005 instance has a default error log path in the format            C:\Program Files\Microsoft SQL Server\MSSQL            \$instance_name\LOG\ERRORLOG,            where <i>instance_name</i> is the Microsoft SQL Server instance name.</p>

Table 50. Names and descriptions of configuration settings in the **Database Server Properties** window (continued)

Parameter name	Description	Mandatory field	Examples
Windows Authentication	<p>Windows Authentication is a Windows account with which the agent services are configured, and is the default configuration option.</p> <p>If you select the <b>Windows Authentication</b> check box, Windows credentials are used for authentication.</p> <p>When the Microsoft SQL Server agent is configured with Windows Authentication, either <b>Local System account</b> or <b>This account</b> is used by the agent services to log on to the Microsoft SQL Server.</p> <ul style="list-style-type: none"> <li>• If the agent services are configured to use <b>Local System account</b> to log on, the agent uses the NT AUTHORITY\SYSTEM user ID to access the Microsoft SQL Server.</li> <li>• If the agent services are configured to use <b>This account</b> to log on, the agent uses the respective user ID to access the Microsoft SQL Server.</li> </ul> <p><b>Remember:</b> If you do not select the <b>Windows Authentication</b> check box, you must specify values for the <b>Login</b> and <b>Password</b> parameters. If you do not specify these parameters and click <b>OK</b> in the <b>Database Server Properties</b> window, an error message is displayed and the agent configuration does not finish.</p> <p><b>Important:</b> If you configure the agent by selecting the <b>Windows Authentication</b> check box and specifying a login ID in the <b>Login</b> field, the agent gives preference to the Windows Authentication.</p>	No	
Support Long Lived Database Connections	<p>Enables or disables long lived database connections. The following data sets do not use long-lived database connections:</p> <ul style="list-style-type: none"> <li>• MS SQL Text</li> <li>• MS SQL Filegroup Detail</li> <li>• MS SQL Server Summary</li> </ul>	No	

Table 50. Names and descriptions of configuration settings in the **Database Server Properties** window (continued)

Parameter name	Description	Mandatory field	Examples
Extended Parms	Disables data collection of any attribute group.	No	<p>For example:</p> <p>To disable the data collection for Table Details data set, enter koqtb1d in the <b>Extended Parms</b> field.</p> <p>To disable the data collection for Database Details data set, enter koqdbd in the <b>Extended Parms</b> field.</p> <p>To disable the data collection for Table Details and Database Details data sets, enter koqtb1d, koqdbd in the <b>Extended Parms</b> field.</p>

Table 50. Names and descriptions of configuration settings in the **Database Server Properties** window (continued)

Parameter name	Description	Mandatory field	Examples
Database	<p>To select the databases for monitoring, specify a value for this parameter. To monitor all the databases that are available on the SQL server instance, select the <b>Monitor All Databases</b> check box in the <b>Databases</b> group area.</p> <p><b>Tip:</b> The <b>Monitor All Databases</b> check box is selected by default.</p> <ul style="list-style-type: none"> <li>To monitor particular databases, select <b>Include</b> from the list, and specify the database names in the field.</li> <li>To exclude particular databases from being monitored, select <b>Exclude</b> from the list, and specify the database names in the field.</li> </ul> <p>Use the field to filter databases that you want to monitor.</p> <p>To specify database filter, you must first select a separator. A separator is a character that separates a database name or database expression from the others.</p> <p>When you are selecting a separator, ensure that database names and database expression do not contain the separator character. You must not use the wildcard characters that are typically used in the T-SQL query (for example, %, _, [ ], ^, -).</p> <p>When you are specifying database filter:</p> <ul style="list-style-type: none"> <li>Database names must start with a separator.</li> <li>Database expression must start with 2 separators.</li> </ul> <p><b>Note:</b> Database expression is a valid expression that can be used in the LIKE part of the T-SQL query. However, you cannot use the T-SQL ESCAPE clause when you are specifying the database expression.</p> <p>The following data sets are affected by database filter:</p> <ul style="list-style-type: none"> <li>Database Detail</li> <li>Database Summary</li> <li>Device Detail</li> <li>Table Detail</li> <li>Table Summary</li> <li>Filegroup Detail</li> <li>Additional Database Detail</li> </ul>	No	<p>Examples of filters:</p> <p>Case 1: % usage</p> <p>Example:</p> <pre>@@%m%</pre> <p>Output: All the databases that have the character <b>m</b> in their names are filtered.</p> <p>Case 2: _ usage</p> <p>Example:</p> <pre>@@_---</pre> <p>Output: All the databases that are of length four characters are filtered.</p> <p>Case 3: [ ] usage</p> <p>Example:</p> <pre>@@[m]---</pre> <p>Output: All the databases of length four characters and whose names start with the character <b>m</b> are filtered.</p> <p>Case 4: [^] usage</p> <p>Example:</p> <pre>@@[^m]%</pre> <p>Output: All the databases (of any length) except the names start with the character <b>m</b> are filtered.</p>

Table 50. Names and descriptions of configuration settings in the **Database Server Properties** window (continued)

Parameter name	Description	Mandatory field	Examples
Database (continued)	<p><b>Remember:</b></p> <ul style="list-style-type: none"> <li>If you do not select the <b>Monitor All Databases</b> check box, you must specify the list of databases for which you want to enable or disable monitoring in the <b>Databases</b> group area. If you do not select the <b>Monitor All Databases</b> check box and the <b>Databases</b> group area is blank, agent configuration cannot be completed.</li> <li>If you select the <b>Monitor All Databases</b> check box and specify the databases to monitor in <b>Databases</b> group area, the setting of <b>Monitor All Databases</b> check box takes precedence. The list of databases that you specify in <b>Databases</b> group area is ignored.</li> </ul>		<p>Case 5: Wrong input</p> <p>Example:</p> <pre>@%m%</pre> <p>Output: None of the databases are filtered.</p> <p>Case 6: Default</p> <p>Example: Field is blank (No query is typed)</p> <p>Output: All the databases are filtered.</p> <p>Case 7: Mixed patterns</p> <p>Example:</p> <pre>@@[m-t]_d%</pre> <p>Output: All the databases name (of any length) start with the characters <b>m, n, o, p, q, r, s, t</b>, followed by any character, with the character <b>d</b> in the third place are filtered.</p>
Day(s) Frequency	Use this feature to define the frequency of collecting data of Table Detail attributes. The values can be from zero to 31.	No	
Weekly Frequency	Use this feature to specify a particular day for collecting data for Table Detail attributes. The values can be from zero to 7.	No	
Monthly Frequency	Use this feature to define the data collection of Table Detail attributes on a particular day of the month. The possible values are 1, 2, 3, and so on.	No	
Collection Start Time	<p>The collection start time can be entered in HH:MM format.</p> <p>The possible values for hours are zero to 23. The default value is zero.</p> <p>The possible values for minutes are from zero to 59. The default value is zero.</p>	No	



Table 50. Names and descriptions of configuration settings in the **Database Server Properties** window (continued)

Parameter name	Description	Mandatory field	Examples
Table Detail Continuous Collection	Use this feature for the continuous background collection of Table Detail data. The <b>Table Detail Continuous Collection</b> check box is selected by default.	No	
Interval Between Two Continuous Collection (in min.)	Specify the time for the interval between two collections in minutes. The minimum interval time is 3 minutes. You can select the <b>Interval Between Two Continuous Collection (in min.)</b> check box or you can use Scheduling to specify continuous collection of the Table Detail data set. If you select the <b>Interval Between Two Continuous Collection (in min.)</b> check box, you must specify the time interval for collection. If you use Scheduling to specify the collection of the Table Detail data set, the minimum time interval is 1 day. The default interval between two continuous collections is 3 minutes.	No	

The agent collects the data at the time interval for which data collection occurs frequently. For example, if you specify all frequencies (daily, weekly, and monthly) for collecting data, the agent starts the data collection according to the following conditions:

- If day(s) frequency  $\leq 7$ , the day(s) frequency settings are selected, and the weekly and monthly frequency settings are ignored.
- If day(s) frequency  $> 7$ , the weekly frequency settings are selected, and the day(s) and monthly frequency settings are ignored.

**Remember:** If the **Table Detail Continuous Collection** check box is selected, the agent collects the data at the interval that is mentioned in the **Interval Between Two Continuous Collection (in min.)** field and ignores the daily, weekly, or monthly frequencies.

## Configuring local environment variables

You can configure local environment variables to change the behavior of the Microsoft SQL Server agent.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, from the **Actions** menu, click **Advanced > Edit Variables**.
3. In the Monitoring Agent for Microsoft SQL Server: **Override Local Variable Settings** window, click **Add**.
4. In the **Add Environment Setting Override** window, enter the variable and the corresponding value.

To view the list of environment variables that you can configure, see [“Local environment variables” on page 390](#).

## Local environment variables

You can change the behavior of the Microsoft SQL Server agent by configuring the local environment variables.

### Variables for checking availability of the SQL Server service

To check the availability of the SQL Server service, use the following environment variables:

- **COLL\_MSSQL\_RETRY\_INTERVAL**: This variable provides the retry interval (in minutes) to check the SQL Server service status. If the value is less than or equal to zero, then the variable takes the default value of 1 minute.
- **COLL\_MSSQL\_RETRY\_CNT**: This variable provides the number of retries that the SQL Server agent makes to check whether the SQL Server service is started or not. If the SQL Server service is not started after the number of retries that are specified in this variable, then collector stops working. If the value of the variable is less than or equal to zero, then the variable takes the default value of 3.

### Variables for monitoring the SQL Server error log file

To monitor the MS SQL Error Event Details data set, use the following environment variables:

- **COLL\_ERRORLOG\_STARTUP\_MAX\_TIME**: This variable provides the time interval (T) for error collection before the agent starts. The default value is 0 minutes. This variable can take the following values:

#### T = 0

The agent starts monitoring the error log file from the time the agent starts or is restarted. The agent does not read the errors that were logged in the error log file before the agent was started.

#### T = 1

The agent monitors the error log file according to the following values that are set for the **COLL\_ERRORLOG\_STARTUP\_MAX\_EVENT\_ROW** variable, which is represented by R:

- If  $R < 0$ , the agent starts monitoring the error log file from the time the agent starts or is restarted.
- If  $R = 1$ , the agent monitors all the errors that are logged in the error log file.
- If  $R > 1$  and the agent is installed for the first time, the agent monitors the error log file until R errors are monitored. If  $R > 1$  and the agent is restarted, the agent monitors all the previously missed R errors.

#### T > 1

The agent monitors all previous errors that were logged up to T minutes from the time that the agent starts or restarts. The agent monitoring also depends on the following values that you set for the **COLL\_ERRORLOG\_STARTUP\_MAX\_EVENT\_ROW** variable:

- If  $R \leq 0$ , the agent starts monitoring the error log file from the time the agent is started or the agent is restarted.
- If  $R = 1$ , the agent monitors the error log file for all the errors that are logged up to T minutes.
- If  $R > 1$ , the agent does not monitor more than R errors that are logged in last T minutes.

- **COLL\_ERRORLOG\_STARTUP\_MAX\_EVENT\_ROW**: This variable provides the maximum number of errors that must be processed when the agent starts. The default value is 0. You can assign following values to this variable:

#### R = 0

The agent starts monitoring the error log file from the time that the agent starts or restarts. The agent does not read errors that were created in the error log file before the agent was started.

#### R = 1

The agent monitors the errors that were logged in the last T minutes from the time that the agent starts or restarts.

#### R > 1

The agent monitors R errors that are logged in the last T minutes.

- **COLL\_ERRORLOG\_MAX\_EVENT\_ROW**: This variable provides the number of error rows. The default value is 50. You can assign following values to this variable:

**X = 0**

The agent does not display the error logs.

**X > 0**

The agent displays the X error rows.

- **COLL\_ERRORLOG\_RECYCLE\_WAIT**: This variable provides the time interval (in seconds) for which the Microsoft SQL Server agent waits before collecting data of the MS SQL Error Event Detail attribute group when the situation on this attribute group is triggered. You can assign a value to this variable in the range of 1 to 30. If the value of this variable is less than zero, then the variable takes the default value of zero (seconds). If the value of this variable is greater than 30, then the variable takes the default value of 30 (seconds).

### Variable for setting the query timeout interval

To set the query timeout interval for the SQL Server agent, use the following environment variables:

- **QUERY\_TIMEOUT**: This environment variable defines the maximum amount of time (in seconds) that the SQL Server agent waits to receive a response for a query that is sent to the SQL Server. The value for this variable must be less than 45 seconds. However, if you set the value for this variable as 0 seconds, the SQL Server agent waits indefinitely to receive a response from the SQL Server. If the SQL Server agent accesses many locked databases, you must assign the value to this variable in the range of 10 - 20 seconds. If the query is not processed within the set timeout interval, the SQL Server agent skips the timed out query and moves to the next query in the queue. The agent does not display data for the query that timed out.
- **QUERY\_THREAD\_TIMEOUT**: This environment variable defines the maximum amount of time (in seconds) that the SQL Server agent waits to receive a response for a query that is sent to the SQL Server. This environment variable is applicable for few attribute groups that uses threaded collection. For example, KOQDBD, KOQTBLD, KOQDEVD, and so on. The value for this variable does not have any limit unlike QUERY\_TIMEOUT variable. Otherwise, it works similar to QUERY\_TIMEOUT variable.

### Variable for viewing information about the enabled jobs

To view the information about enabled jobs in the MS SQL Job Detail data set, use the

**COLL\_JOB\_DISABLED** environment variable. If you set the value of this variable as 1, the Microsoft SQL Server agent does not display information about disabled jobs. If you do not specify this variable, you can view information that is about enabled and disabled jobs.

### Variable for limiting the rows in the MS SQL Filegroup Detail data set

To limit the number of rows that the collector service fetches for the MS SQL Filegroup Detail data set, use the **COLL\_KOQFGRPD\_MAX\_ROW** environment variable. This environment variable defines the maximum number of rows that the collector service fetches for the Filegroup Detail data set. If you do not specify a value for this variable, the collector service fetches 10,000 rows for the Filegroup Detail data set. Use this environment variable to modify the default limit of maximum rows in the `koqcoll.ct1` file. Complete the following steps to modify the default limit:

1. Specify the maximum number of rows for KOQFGRPD in the `koqcoll.ct1` file.
2. Add the **COLL\_KOQFGRPD\_MAX\_ROW** environment variable, and ensure that the value of this variable is the same as the value that you have specified in the `koqcoll.ct1` file.

If the value in the `koqcoll.ct1` file is less than the value that is specified in the

**COLL\_KOQFGRPD\_MAX\_ROW** environment variable, the value in the `koqcoll.ct1` file is treated as the value for the maximum number of rows.

If the value in the `koqcoll.ct1` file is greater than the value that is specified in the

**COLL\_KOQFGRPD\_MAX\_ROW** environment variable, the value in the **COLL\_KOQFGRPD\_MAX\_ROW** environment variable is treated as the value for the maximum number of rows.

### Variables for enhancing the collection for the MS SQL Filegroup Detail data set

Use the **COLL\_DBD\_FRENAME\_RETRY\_CNT** variable to specify the number of attempts that can be made to move the %COLL\_HOME%\_tmp\_%COLL\_VERSION%\_%COLL\_SERVERID%\_%COLL\_SERVERID%\_%FGRP\_TEMP file to the %COLL\_HOME%\_tmp\_%COLL\_VERSION%\_%COLL\_SERVERID%\_%COLL\_SERVERID%\_%FGRP\_PREV file.

If you do not specify a value for this variable, the Microsoft SQL Server agent makes 3 attempts to move the file.

### Variable for limiting the rows in the MS SQL Device Detail data set

To limit the number of rows that the collector service fetches for the MS SQL Device Detail data set, use the **COLL\_KOQDEVD\_MAX\_ROW** environment variable. This environment variable defines the maximum number of rows that the collector service fetches for the Device Detail data set. If you do not specify a value for this variable, the collector service fetches 10,000 rows for the Device Detail data set. Use this environment variable to modify the default limit of maximum rows in the koqcoll.ctl file. Complete the following steps to modify the default limit:

1. Specify the maximum number of rows for KOQDEVD in the koqcoll.ctl file.
2. Add the **COLL\_KOQDEVD\_MAX\_ROW** environment variable, and ensure that the value of this variable is the same as the value that you have specified in the koqcoll.ctl file.

If the value in the koqcoll.ctl file is less than the value that is specified in the **COLL\_KOQDEVD\_MAX\_ROW** environment variable, the value in the koqcoll.ctl file is treated as the value for the maximum number of rows.

If the value in the koqcoll.ctl file is greater than the value that is specified in the **COLL\_KOQDEVD\_MAX\_ROW** environment variable, the value in the **COLL\_KOQDEVD\_MAX\_ROW** environment variable is treated as the value for the maximum number of rows.

### Variables for enhancing the collection for the MS SQL Device Detail data set

To enhance the MS SQL Device Detail data set collection, use the following environment variables:

- **COLL\_KOQDEVD\_INTERVAL**: This environment variable enables you to specify a time interval (in minutes) between two consecutive collections of the MS SQL Device Detail data set.

**Note:** By default, the data collection for the Device Detail data set is demand based. Use the **COLL\_KOQDEVD\_INTERVAL** variable to start a thread based collection for the Device Detail data set and to set the time interval between two threaded collections.

- **COLL\_DBD\_FRENAME\_RETRY\_CNT**: Use this environment variable to specify the number of attempts that can be made to move the %COLL\_HOME%\_tmp\_%COLL\_VERSION%\_%COLL\_SERVERID%\_%COLL\_SERVERID%\_DEVD\_TEMP file to the %COLL\_HOME%\_tmp\_%COLL\_VERSION%\_%COLL\_SERVERID%\_%COLL\_SERVERID%\_DEVD\_PREV file.

If you do not specify a value for this variable, the Microsoft SQL Server agent makes 1 attempt to move the file.

### Variables for enhancing the collection for the MS SQL Database Detail data set

To enhance the MS SQL Database Detail data set collection, use the following environment variables:

- **COLL\_KOQDBD\_INTERVAL**: Use this environment variable to specify a time interval (in minutes) between two consecutive thread-based collections of the MS SQL Database Detail data set. If you do not specify a value for this variable or the specified time interval is less than 3 minutes, then the Microsoft SQL Server agent defaults to 3 minutes interval. In case, the collection is taking more time or the data is frequently seen as NOT\_COLLECTED, then you can check the collection time by referring to the Database Detail Collection completed in %d seconds log and set the variable value to a value that is greater than the collection time specified in the log.
- **COLL\_DBD\_FRENAME\_RETRY\_CNT**: Use this environment variable to specify the number of attempts that can be made to move the %COLL\_HOME%\_tmp\_%COLL\_VERSION%\_%COLL\_SERVERID%\_

%COLL\_SERVERID%\_DBD\_TEMP file to the %COLL\_HOME%\_tmp\_%COLL\_VERSION%\_%COLL\_SERVERID%\_%COLL\_SERVERID%\_DBD\_PREV file.

If you do not specify a value for this variable, the Microsoft SQL Server agent makes 1 attempt to move the file.

### Variables for enhancing the collection for the MS SQL Audit Details data set

To enhance the MS SQL Audit Details data set collection, use the following environment variables:

- **COLL\_AUDIT\_TYPE**: Use this variable to enable or disable the monitoring of specific logs. The default value of the variable is [AL][FL][SL]. By default, the agent monitors all three types of logs that include the application logs, audit files, and the security logs. The value of the variable includes two character code for each log type:

- [AL] for application logs
- [FL] for audit files
- [SL] for security logs

You can change the value of the variable to disable the monitoring of specific log type. For example, if you specify the value of the variable as [AL][SL] the audit files are not monitored. If no value is specified for the variable, audit details not monitored.

- **COLL\_AUDIT\_DURATION**: Use this variable to report the audit events that occurred during the time interval that you specify in this variable. For example, if you set this variable to 7, the audit events that occurred only in last 7 hours are reported by the Audit Details data set. The default value of the **COLL\_AUDIT\_DURATION** variable is 24 hours.
- **COLL\_AUDIT\_COLLECTION\_INTERVAL**: The threaded collection in the Audit Details data set provides specifications of all the database that are present on the SQL server instance. Use this variable to set the interval for this threaded collection. For example, if you set this variable to 7, a fresh set of database specifications is extracted from the SQL server instance after every 7 hours. The default value of the **COLL\_AUDIT\_COLLECTION\_INTERVAL** variable is 24.

### Variable for enhancing the collection for the MS SQL Process Detail data set

To enhance the MS SQL Process Detail data set collection, use the **COLL\_PROC\_BLOCK\_INTERVAL** variable with the following values:

- If **COLL\_PROC\_BLOCK\_INTERVAL** = 0, the collection for the Blocking Process Duration attribute, and the Blocking Resource Duration attribute is disabled.
- If **COLL\_PROC\_BLOCK\_INTERVAL** = x, the interval between the two consecutive data collections for the Blocking Process Duration and the Blocking Resource Duration attributes is x minutes.

If the **COLL\_PROC\_BLOCK\_INTERVAL** variable is not set in the CANDLE\_HOME directory, the interval between the two consecutive data collections is three minutes.

### Variable for excluding the locked objects from the data collection

If the queries that are sent for the Database Detail, Filegroup Details, Database Mirroring, and Device Detail workspaces take long to execute, use the **COLL\_DBCC\_NO\_LOCK** variable to run a query with the value WITH (NOLOCK). This variable causes the query not to wait in the queue when an object on which the query is run is locked.

### Variable for setting the sorting criteria for the rows returned by the Table Details data set

The rows that are returned by the Table Details data set are sorted in a descending order depending on the value that is set for the **COLL\_TBLD\_SORTBY** variable. The default value for the **COLL\_TBLD\_SORTBY** variable is FRAG (fragmentation percent). The valid values are: ROWS (number of rows in a tables), SPACE (space used by the table), and OPTSAGE (the optimizer statistics age of the table).

## Variable for enhancing the collection for the MS SQL Problem Detail and Problem Summary data sets

- **COLL\_ALERT\_SEV**: Use this variable to set the severity level of the error messages that are displayed in the Problem Detail and Problem Summary data sets. Error messages, which have a severity level that is equal to or greater than the value mentioned in this variable, are displayed in the Problem Detail and Problem Summary data sets. For example, if you set the value of this variable to 10, the error messages with severity level 10 or greater are displayed in the Problem Detail and Problem Summary data sets. If you do not specify a value for this variable, the error messages, which have a severity level that is equal to or greater than 17, are displayed in the Problem Detail and Problem Summary data sets.
- **COLL\_SINCE\_ERRORLOG\_RECY**: Use this variable to monitor only the high severity errors in the current ERRORLOG file. If you do not specify a value for this variable, the value of the variable is 0, which means that for collecting the data, the Problem Summary data set also considers the high severity errors that are read from the previous ERRORLOG file. To monitor only the high severity errors in the current ERRORLOG file, set the value of this variable to 1.

## Variables for setting the timeout interval

To set the timeout interval for the Microsoft SQL Server agent, you can use the following environment variables:

- **WAIT\_TIMEOUT**: Use this variable to set the wait timeout interval for the Microsoft SQL Server agent. If any data set takes more than 45 seconds to collect data, then the agent might hang or situations might be incorrectly triggered. Check the log for the data sets that take more than 45 seconds to collect the data, and use the **WAIT\_TIMEOUT** variable to increase the wait time between the agent process and the collector process.
- **COLL\_DB\_TIMEOUT**: Use this variable to define the wait interval (in seconds) for any request such as running a query on the existing SQL server connection to complete before returning to the application. If you set this value to 0, then there is no timeout. If you do not specify a value for this variable, the agent waits 15 seconds before returning to the application.

## Variables for setting the properties of the collector log files

To set the properties of the collector log files, you can use the following environment variables:

- **COLL\_WRAPLINES**: Use this variable to specify the maximum number of lines in a col.out file. The default value of this variable is 90,000 lines (about 2 MB).
- **COLL\_NUMOUTBAK**: Use this variable to specify the number of backup copies of the collector log files that you want to create. By default, five backup copies of the collector log file are created. The backup file is named \*.out. When this backup file is full, the file is renamed to \*.ou1 and the latest logs are written in the \*.out file. In this manner, for five backup files, the oldest logs are available in the \*.ou5 file and the latest logs are available in the \*.out file.

You can create more than five backup copies of the collector log files by specifying one of the following values in the **COLL\_NUMOUTBAK** variable:

- For less than 10 backup files, specify the number of backup files that you want to create in the **COLL\_NUMOUTBAK** variable. For example, if you specify 9 in the **COLL\_NUMOUTBAK** variable, nine backup files will be created.
- For more than 9 and less than 1000 backup files, in the **COLL\_NUMOUTBAK** variable, specify the number of backup files preceded by a hyphen. For example, if you specify -352 in the **COLL\_NUMOUTBAK** variable, three hundred and fifty-two backup files will be created.
- **COLL\_DEBUG**: Use this variable to enable full tracing of the collector by setting the value of this variable to dddddddddd (10 times"d").

## Variable for deleting the temporary files

**COLL\_TMPFILE\_DEL\_INTERVAL**: Use this variable to specify the interval (in minutes) after which the KOQ\_<timestamp> temporary files should be deleted. If you do not specify a value for this variable, the value of the variable is 0, which means that the temporary files must be deleted immediately.

### Variable for changing driver used by the MS SQL Server agent

To change the driver that is used by the Microsoft SQL Server agent, use the **KOQ\_ODBC\_DRIVER** environment variable. This variable specifies the driver that the Microsoft SQL Server agent uses to connect to the SQL Server. If you do not specify a value for this variable, then agent uses the ODBC SQL Server Driver as a default driver.

**Note:** When you specify the Microsoft SQL Server driver, ensure that the driver name is correct and the driver is listed under the drivers' option in data source (ODBC).

### Variable for connecting to an AlwaysOn enabled SQL Server database

**KOQ\_APPLICATION\_INTENT:** Use this variable to specify the connection option while connecting to SQL Server.

**KOQ\_APPLICATION\_INTENT** option details:

- **Readonly:** Connection is opened with **ApplicationIntent** as *readonly*.
- **Readwrite:** Connection is opened with **ApplicationIntent** as *readwrite*.  
When it is set to Readwrite, Microsoft SQL Server agent would not perform any write operations with the connection.

If this variable is not set, the connection is established without **ApplicationIntent** property.

**Note:** The driver is specified by the environment variable **KOQ\_ODBC\_DRIVER**. If this variable is not set, then the default SQL Server driver is used.

If the driver doesn't support **ApplicationIntent**, the connection is opened without **ApplicationIntent** property.

## Configuring the agent by using the silent response file

You can use the silent response file for configuring the agent. You can also configure multiple instances of the agent by using the silent response file.

### Before you begin

To configure multiple instances of the agent, ensure that the configuration details of all the agent instances are specified in the silent response file.

### About this task

Run the configuration script to change the configuration settings. You can edit the silent response file before you run the configuration script.

### Procedure

To configure the agent, complete the following steps:

1. Open the `mssql_silent_config.txt` file that is at `install_dir\samples`, and specify values for all mandatory parameters.

You can also modify the default values of other parameters.

2. Open the command prompt, and enter the following command:

```
install_dir\BIN\mssql-agent.bat config install_dir\samples  
\mssql_silent_config.txt
```

3. Start the agent.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Running the agent in a cluster environment

You can configure the Microsoft SQL Server agent in a cluster environment. Multiple instances of the Microsoft SQL Server and the Microsoft SQL Server agent can run on a single node.

After you install and configure the Microsoft SQL Server agent, complete the following tasks to run the agent in a cluster environment:

- Add environment variables
- Change the startup type of the agent service and the collector service
- Add the agent and the collector to the cluster environment

You can set up a cluster environment for the following versions of the Microsoft SQL Server:

- Microsoft SQL Server 2005
- Microsoft SQL Server 2008
- Microsoft SQL Server 2008 R2
- Microsoft SQL Server 2012
- Microsoft SQL Server 2014
- Microsoft SQL Server 2016

**Important:** On Windows systems, the agent must be installed in the same directory where the OS agent is installed. Install the agent on the nodes system disk of each cluster node.

### Adding environment variables

You must configure the environment variables that are used by the agents that are installed on each cluster node.

### About this task

You must specify values for the following environment variables:

- **CTIRA\_HOSTNAME:** This variable is used to configure each instance of the Microsoft SQL Server agent. The value of this variable is limited to 31 characters and is common for all monitoring agents. Set the value of this variable to the cluster name. Users can navigate to all the monitoring agents of that cluster in the Cloud App Management console.
- **CTIRA\_NODETYPE:** This variable is used to identify the agent. By default, the value of this variable is set to **MSS** for the Microsoft SQL Server agent.
- **CTIRA\_SUBSYSTEMID:** This variable is used to distinguish the multiple instances of the Microsoft SQL Server agent. By default, the value of this variable is set to **Microsoft SQL Virtual Server** for the Microsoft SQL Server agent.
- **COLL\_HOME:** This variable is used to collect data and store log files for attribute groups that use configuration files at a shared location. Set the value of the variable to *X:\shared-location*, where *X* is a shared drive that is accessible to the cluster nodes. For example, set the value for the **COLL\_HOME** when you define the configuration settings for the MS SQL Table Detail attribute group or MS SQL Error Event Details attribute group.
- **CTIRA\_HIST\_DIR:** This variable specifies the path to the shared disk directory. If history for the Microsoft SQL Server agent is configured to be stored at the monitoring agent, each instance of the agent must be configured with a common **CTIRA\_HIST\_DIR** variable that refers to the shared disk directory.

**Remember:** If history is stored at the Cloud App Management server, you need not specify a value for the **CTIRA\_HIST\_DIR** variable. Storing history at the Cloud App Management server increases the load on that server.

To add these variables, see the steps that are described in [“Configuring local environment variables” on page 389](#).



### What to do next

Change the startup type of the agent service and the collector service to **Manual** by completing the steps that are described in [“Changing the startup type of the agent service and the collector service”](#) on page 397.

### Changing the startup type of the agent service and the collector service

By default the startup type of the agent service and the collector service is **Automatic**. Change the startup type of the agent service and the collector service to **Manual** so that the cluster resource can control the starting and stopping of the monitoring agent

#### Procedure

To change the startup type of the agent service, complete the following steps:

1. Click **Start > Run**, type the command `services.msc`, and then click **OK**.
2. Right-click the agent and click **Properties**.
3. In the **Monitoring Agent for Microsoft SQL Server Properties** window, from the **Startup type** list, select **Manual**, click **Apply**, and then **OK**.

### What to do next

- Use the same procedure to change the startup type of the collector service to **Manual**.
- Add the agent and the collector to the cluster environment by completing the steps that are described in [“Adding the agent and collector to the cluster environment ”](#) on page 397.

### Adding the agent and collector to the cluster environment

You must add the agent and the collector to the cluster environment.

#### Procedure

1. Click **Start > Control Panel > Administrative Tools > Failover Cluster Management**.
2. Expand **Failover Cluster Management**.
3. Expand **Services and Applications** and right-click the SQL instance that you want to configure.
4. Click **Add a resource > Generic Service**. The New Resource Wizard opens.
5. On the Select Service page, select the service name, and then click **Next**.

Examples of Windows Services names:

- Monitoring Agent for Microsoft SQL Server: `SQLTEST#INSTANCE1`
  - Monitoring Agent for Microsoft SQL Server: Collector `SQLTEST#INSTANCE1`
  - Monitoring Agent for Microsoft SQL Server: `SQLTEST2#INSTANCE2`
  - Monitoring Agent for Microsoft SQL Server: Collector `SQLTEST2#INSTANCE2`
6. On the Confirmation page, check the details, and then click **Next**.
  7. On the Summary page, click **Finish**. The Microsoft SQL Server agent is now added.

**Remember:** Use the same steps to add the collector to the cluster environment.

8. To bring the agent online, right-click the agent, and click **Bring this resource online**.
9. To bring the collector online, right-click the collector, and click **Bring this resource online**.

### Results

The Microsoft SQL Server agent is now running in a cluster environment.

**Remember:** If you want to configure the agent again, you must first take the agent and the collector offline, or edit the agent variables on the node where the agent and collector run. When you complete the agent configuration, bring the agent and the collector back online.

## Configuring the agent by using the cluster utility

You can use the cluster utility to add multiple Microsoft SQL Server agent instances to a cluster group in a cluster environment.

The cluster utility automatically adds the agent service and the collector service of each Microsoft SQL Server agent instance as a generic service resource to the cluster group. You can use the cluster utility to complete the following tasks:

- [Adding an SQL Server agent instance to the cluster](#)
- [Updating an existing SQL Server agent instance in a cluster](#)
- [Removing an SQL Server agent instance from a cluster](#)

### Prerequisites for using the cluster utility

You must ensure that your system environment meets the prerequisites for running the cluster utility.

Ensure that the following prerequisites are met:

- Run the cluster utility on a computer that has at least one group in the cluster environment.
- Start the remote registry service for all nodes in the cluster.
- You must have the cluster manager authorization to access the cluster utility.
- The service name of agent and collector must be same on all cluster node.

For example, if the agent service name is Monitoring Agent for Microsoft SQL Server: SQLTEST#INSTANCE1 and the collector name is Monitoring Agent for Microsoft SQL Server: Collector SQLTEST#INSTANCE1 then the same service name must be present on all nodes of cluster.

### Adding an Microsoft SQL Server agent instance to the cluster

You can use the cluster utility to add an Microsoft SQL Server agent instance to a cluster group in a cluster environment.

#### Procedure

1. To run the utility, complete one of the following steps:
  - For a 64-bit agent, go to the *candle\_home*\TMAITM6\_x64 directory.
  - For a 32-bit agent, go to the *candle\_home*\TMAITM6 directory.
2. To run the Cluster Utility, double-click the *KoqClusterUtility.exe*.
3. In the **SQL Server Agent Instances Available** area, select a Microsoft SQL Server agent instance, and click **Add**.
4. In the **Select cluster group name** window, select a cluster group.

The cluster group that you select must be the SQL Server instance that is monitored by the Microsoft SQL Server agent.
5. In the **Select Path for Shared Location** window, navigate to the path where the agent and collector logs are stored.

If you do not select the path, by default, the *CANDLEHOME/TMAITM6(\_x64)/logs* location is selected for storing the agent and collector logs.
6. To add the Microsoft SQL Server agent instance to the cluster environment, click **OK**.

The activity logs of the cluster utility are displayed in the **History** pane.

### Updating an existing Microsoft SQL Server agent instance in a cluster

You can use the cluster utility to update the location where the agent and collector logs are stored for an SQL Server instance in a cluster.

#### Procedure

1. To update an existing Microsoft SQL Server agent instance, open the **Cluster Utility** window.

2. In the **SQL Server Agent Instances Configured** area, select a Microsoft SQL Server agent instance, and click **Update**.
3. In the **Set Path for Shared Location** window, navigate to the path where the agent and collector logs are stored.

If you do not select the path, the agent and collector logs are stored at the location that was set while adding the Microsoft SQL Server agent instance in a cluster.

4. Click **OK**.

The activity logs of the cluster utility are displayed in the **History** pane.

### Removing a Microsoft SQL Server agent instance from a cluster

You can use the cluster utility to remove a Microsoft SQL Server agent instance from a cluster group.

#### Procedure

1. Open the **Cluster Utility** window.
2. In the **SQL Server Agent Instances Configured** area, select a Microsoft SQL Server agent Instance, and click **Remove**.
3. In the **Please Confirm Action** dialog box, click **Yes** to delete the Microsoft SQL Server agent instance from the cluster.

The activity logs of the cluster utility are displayed in the **History** pane.

## Configuring multiple collations for ERRORLOG file

The Microsoft SQL Server agent supports multiple collations in ERRORLOG file. You can configure the agent to parse multiple collations in the ERRORLOG file for **Problem Detail** attribute group. Multiple collations in ERRORLOG file are not applicable for **Error Event Detail** attribute group.

### Before you begin

To configure multiple collations of the agent, ensure that the agent is installed.

### About this task

The default collation is English. For other languages of SQL Server, the agent parses the ERRORLOG file based on the collations in the configuration file `koqErrConfig.ini`. So you must add the collations that are in used in `koqErrConfig.ini` file.

#### Procedure

To configure multiple collations for the agent, complete the following steps:

1. Go to the agent directory `agent_directory`, where:
  - For 64-bits agent, the `agent_directory` is `Agent_home\TMAITM6_x64`.
  - For 32-bits agent, the `agent_directory` is `Agent_home\TMAITM6`.

The `Agent_home` is the agent installation directory.

2. Open the configuration file `koqErrConfig.ini` with your editor.
3. Add the new collations.

For example, to enable French collation, add the following collation settings in **name-value** pair format in `koqErrConfig.ini`.

```
[French]
Error = Erreur :
Severity = Gravité :
State = État :
```

**Note:** The sample list of collations is available in `agent_directory\koqErrConfigSample.ini`, where:

- For 64-bits agent, the *agent\_directory* is *Agent\_home\TMAITM6\_x64*.
- For 32-bits agent, the *agent\_directory* is *Agent\_home\TMAITM6*.

The *Agent\_home* is the agent installation directory.

If the target collation is not available in *koqErrConfigSample.ini*, you can determine the collation keyword values from the ERRORLOG file.

Adhere to the following collation format when configure the collation settings in *koqErrConfig.ini*.

```
[Section_name]
Error = Error_value
Severity = Severity_value
State = State_value
```

- The *Section\_name* is the SQL Server collation name. Ensure that the collation name is enclosed with an open bracket "[" and a closed bracket "]".
- The *Error\_value* is the error keyword found in ERRORLOG file of your target collation.
- The *Severity\_value* is the severity keyword found in ERRORLOG file of your target collation.
- The *State\_value* is the state keyword found in ERRORLOG file of your target collation.

**Important:** The keyword values must be the same as the keyword values found in the ERRORLOG file, including any special characters.

#### 4. Save the configuration file *koqErrConfig.ini*.

Agent restart is not required.

If the configuration file *koqErrConfig.ini* is not available or the configuration file *koqErrConfig.ini* is empty, the ERRORLOG file shows the default collation as English error message. The error message severity level is more than the default severity level, if any.

If the configuration file *koqErrConfig.ini* is configured correctly, the ERRORLOG file shows the error messages with severity level more than the default severity level, if any.

The default severity level is 17.



**Attention:** Before agent upgrade, you must make a copy of the *koqErrConfig.ini* file. It is not preserved during agent upgrade.

#### What to do next

Check the **Errorlog Alert** widget or the **Problem Detail** attribute group on the Cloud App Management console as the result of the collation settings.

## Configuring historical job count for monitoring

You can configure the maximum historical job count for monitoring to display the **Success count** and **Non-success count** in the **Job details** widget. The default value is **100**.

#### Procedure

1. Launch the **SQL Server Management Studio**.
2. Right-click the **SQL Server Agent** and select the **Properties**.
3. On **SQL Server Agent Properties** window, select the **History** page.
4. In the **Maximum job history rows per job** field, enter the rows count and then click OK.

#### Results

The **Job details** widget displays the **Success count** and **Non-success count** for the selected job.

## Configuring MongoDB monitoring

The Monitoring Agent for MongoDB requires an instance name. You must manually configure and start the agent instance. The MongoDB agent supports local as well as remote monitoring.

### Before you begin

- Ensure that the user, who configures the MongoDB agent, has the required roles to collect data for all attributes.
  - To configure the agent on the MongoDB database version 2.4 and version 2.6, the clusterAdmin, readAnyDatabase, and dbAdminAnyDatabase roles must be assigned to the user
  - To configure the agent on the MongoDB database version 3.x and 4.x, the clusterMonitor, readAnyDatabase, and dbAdminAnyDatabase roles must be assigned to the user

To know about the attribute groups for which these user roles are required, see [Table 51 on page 401](#).

- Use an existing user or create a user in the admin database.

**Important:** Before you create a user and grant the required roles to the user, make sure to connect to the MongoDB database and change the database to admin database. If the mongod or mongos process is running in the authentication mode, enter the required credentials to connect to MongoDB database.

1. Run the following command to connect to the MongoDB database:

```
mongo IP:port
```

Where

- *IP* is the IP address of the mongod or mongos process
- *port* is the port number of the mongod or mongos process

2. Change the database to the admin database:

**use admin**

3. Run one of the following commands to add a user in the MongoDB admin database and assign the required roles to the user:

- For the MongoDB database version 2.4, run the following command:

```
db.addUser({ user: "username", pwd: "password", roles: [ 'clusterAdmin', 'readAnyDatabase', 'dbAdminAnyDatabase' ] })
```

- For the MongoDB database version 2.6, run the following command:

```
db.createUser({user: "username", pwd: "password", roles: [ 'clusterAdmin', 'readAnyDatabase', 'dbAdminAnyDatabase' ] })
```

- For the MongoDB database version 3.x and 4.x, run the following command:

```
db.createUser({user: "username", pwd: "password", roles: [ 'clusterMonitor', 'readAnyDatabase', 'dbAdminAnyDatabase' ] })
```

4. Run the following command to verify that the user is added to the admin database:

```
db.auth("username", "password")
```

Return code 1 indicates that the user is added, whereas the return code 0 indicates that the user addition failed.

The following table contains information about the user roles and the attributes for which these user roles are required:

Roles	MongoDB database version	Attribute groups
dbAdminAnyDatabase	2.x, 3.x, 4.x	Response Times

Table 51. Attributes groups and their required user roles (continued)

Roles	MongoDB database version	Attribute groups
readAnyDatabase	2.x, 3.x, 4.x	<ul style="list-style-type: none"> <li>• Mongod Listing</li> <li>• General Shard Information</li> <li>• Collection Storage</li> <li>• Database Names</li> <li>• Shard Details</li> <li>• Collection Storage Details</li> </ul>
clusterAdmin	2.x, 3.x, 4.x	<ul style="list-style-type: none"> <li>• Mongo Instance Information</li> <li>• Mongo Inst IO Info</li> <li>• MII Copy For APMUI One</li> <li>• MII Copy For APMUI Two</li> <li>• Mongo Inst DB Lock</li> <li>• Locks</li> <li>• MongoDB Locks</li> <li>• WiredTiger Details</li> <li>• MMAPv1 Details</li> </ul>
clusterMonitor	2.x, 3.x, 4.x	<ul style="list-style-type: none"> <li>• Mongo Instance Information</li> <li>• Mongo Inst IO Info</li> <li>• MII Copy For APMUI One</li> <li>• MII Copy For APMUI Two</li> <li>• Mongo Inst DB Lock</li> <li>• Locks</li> <li>• MongoDB Locks</li> <li>• WiredTiger Details</li> <li>• MMAPv1 Details</li> </ul>

- For remote monitoring of the MongoDB server, see the two prerequisites
  1. Since MongoDB agent requires mongo shell to collect information remotely from the MongoDB server, the system on which MongoDB agent is installed and configured must have an instance of MongoDB server. The mongo shell of the MongoDB server on the agent machine is used to connect to the remote MongoDB server for monitoring.
  2. In /etc/hosts file of the system that hosts the agent, there is an entry of the remote machine.

### About this task

The managed system name includes the instance name that you specify. For example, you can specify the instance name as *instance\_name: host\_name: pc*, where *pc* is the two character product code of your agent. The managed system name can contain up to 32 characters. The instance name can contain up to 28 characters, excluding the length of your host name. For example, if you specify Mongo2 as your instance name, your managed system name is Mongo2: hostname: KJ.

**Important:** If you specify a long instance name, the managed system name is truncated and the agent code is not completely displayed.

The product version and the agent version often differ. The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see

“Using agent commands” on page 226. For detailed information about the agent version list and what's new for each version, see the [“Change history”](#) on page 52.

### Remember:

- For the agent to successfully collect data, start the agent with the super (root) user, or use the same user ID to start the agent and the mongod process.
- In an environment where MongoDB runs as a cluster, ensure that you install the agent on the same system where the router process is running. Configure the agent on the same system with the IP address and port number of that system and the setup **TYPE** as 1.
- In an environment where MongoDB runs as a cluster in authentication mode, ensure that you add the same user ID with the required rights on all the shards in the cluster.

You can configure the agent by using the default settings, by editing the silent response file, or by responding to prompts.

## Configuring the agent with default settings

For a typical environment, use default settings to configure the agent. When default settings are used for the agent configuration, the agent does not run in the authentication mode.

### Procedure

1. Run the following command:

```
install_dir/bin/mongodb-agent.sh config instance_name install_dir/samples/  
mongodb_silent_config.txt
```

Where

- *instance\_name* is the name that you specify for the unique application instance.
- *install\_dir* is the installation directory of the MongoDB agent.

The default installation directory is `/opt/ibm/apm/agent`.

2. Run the following command to start the agent:

```
install_dir/bin/mongodb-agent.sh start instance_name
```

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

If you need help with troubleshooting, see the [IBM Cloud App Management & IBM Cloud Application Performance Management](#) on developerWorks.

## Configuring the agent by responding to prompts

To configure the agent with custom settings, you can specify values for the configuration parameters when prompted while the script is being run.

### Procedure

1. Run the following command:

```
install_dir/bin/mongodb-agent.sh config instance_name
```

Where

- *instance\_name* is the name that you specify for the instance.
- *install\_dir* is the installation directory of the MongoDB agent.

2. When you are prompted to provide a value for the **TYPE** parameter, press Enter to accept the default value, or specify one of the following values, and then press Enter:

- 1 for a cluster

- 2 for a replication set
- 3 for a stand-alone instance

By default, the agent monitors a cluster.

3. When you are prompted to provide a value for the **PORT** parameter, press `Enter` to accept the default value, or specify the port number of the router for a MongoDB cluster or a `mongod` instance of the MongoDB replication set that is being monitored, and then press `Enter`.

**Remember:** If you do not specify any port number, the agent automatically discovers the port number of the appropriate MongoDB process that is active on the default interface. If no MongoDB process is active on the default interface, then the agent selects the port number of the appropriate MongoDB process that is active on the secondary interface.

4. When you are prompted to provide a value for the **HOST** parameter, press `Enter` to accept the default value, or specify the IP address of the MongoDB host system, and then press `Enter`.

**Remember:** If you do not specify any IP address, the agent automatically detects the IP address of the appropriate MongoDB process that is active on the default interface. If no MongoDB process is active on the default interface, then the agent detects the IP address of the appropriate MongoDB process that is active on the secondary interface.

5. When you are prompted to provide a value for the **AUTHENTICATION** parameter, press `Enter` to accept the default value, or specify whether the agent is running in the authentication mode.

The default value is `NO`, which indicates that the agent is not running in the authentication mode. Specify `YES` to indicate that `mongoDB` is running in the authentication mode.

**Remember:** When the MongoDB database is running in the authentication mode, the MongoDB agent or any MongoDB client cannot connect to the MongoDB database without credentials. To connect to the database that runs in the authentication mode, specify `YES` for the **AUTHENTICATION** parameter.

If you specify `YES`, complete the following steps:

- a) For the **User Name** parameter, specify a user name for the router or the `mongod` instance. Ensure that minimum roles are assigned to the user. For information about user roles, see [Table 51 on page 401](#).
  - b) For the **Password** parameter, specify the password.
6. Run the following command to start the agent:
 

```
install_dir/bin/mongodb-agent.sh start instance_name
```

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

If you need help with troubleshooting, see the [IBM Cloud App Management & IBM Cloud Application Performance Management on developerWorks](#).

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters with default values defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters, and configure the agent.

### Before you begin

To run the MongoDB database in the authentication mode, ensure that you configure the agent with a user who has the `clusterAdmin`, `readAnyDatabase`, and `dbAdminAnyDatabase` roles on the MongoDB database.

### Procedure

1. In a text editor, open the silent response file that is available at the following path:



`install_dir/samples/mongodb_silent_config.txt`.

2. For the **TYPE** parameter, enter one of the following values:

- 1 for a cluster
- 2 for a replication set
- 3 for a stand-alone instance

By default, the agent monitors a cluster.

3. For the **PORT** parameter, specify the port number of the router for a MongoDB cluster or a mongod instance of the MongoDB replication set that is being monitored.

**Remember:** If you do not specify any port number, the agent automatically discovers the port number of the appropriate MongoDB process that is active on the default interface. If no MongoDB process is active on the default interface, then the agent selects the port number of the appropriate MongoDB process that is active on the secondary interface.

4. For the **HOST** parameter, specify the IP address of the MongoDB host system.

**Remember:** If you do not specify any IP address, the agent automatically detects the IP address of the appropriate MongoDB process that is active on the default interface. If no MongoDB process is active on the default interface, then the agent detects the IP address of the appropriate MongoDB process that is active on the secondary interface.

5. For the **AUTHENTICATION** parameter, specify YES to indicate that mongoDB is running in the authentication mode. The default value is NO, which indicates that the agent is not running in the authentication mode.

**Remember:** When the MongoDB database is running in the authentication mode, the MongoDB agent or any MongoDB client cannot connect to the MongoDB database without credentials. To connect to the database that runs in the authentication mode, specify YES for the **AUTHENTICATION** parameter.

If you specify YES, complete the following steps:

- a) For the **User Name** parameter, specify a user name for the router or the mongod instance. Ensure that minimum roles are assigned to the user. For information about user roles, see [Table 51 on page 401](#).
  - b) For the **Password** parameter, specify the password.
6. Save and close the `mongodb_silent_config.txt` file, and run the following command:
- ```
install_dir/bin/mongodb-agent.sh config instance_name install_dir/samples/mongodb_silent_config.txt
```

Where

- `instance_name` is the name that you specify for the instance.
- `install_dir` is the installation directory of the MongoDB agent.

7. Run the following command to start the agent:

```
install_dir/bin/mongodb-agent.sh start instance_name
```

**Important:** If you upgrade the agent to V1.0.0.9 or later and want to run the agent in the authentication mode, then you must configure the agent again to provide a user name and a password. For collecting data, you must stop and restart the agent after configuration.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

If you need help with troubleshooting, see the [IBM Cloud App Management & IBM Cloud Application Performance Management on developerWorks](#).

## Configuring MySQL monitoring

The Monitoring Agent for MySQL requires an instance name and the MySQL server user credentials. You can change the configuration settings after you create the first agent instance.

### Before you begin

- Ensure that a user is created in the MySQL database for running the agent. The user does not require any specific privileges on the MySQL database that is being monitored.
- The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history” on page 52](#).

### About this task

The managed system name includes the instance name that you specify, for example, `instance_name:host_name:pc`, where `pc` is the two character product code. The managed system name can contain up to 32 characters. The instance name that you specify can contain up to 28 characters, excluding the length of your host name. For example, if you specify MySQL2 as your instance name, your managed system name is MySQL2:hostname:SE.

**Important:** If you specify a long instance name, the managed system name is truncated and the agent code is not completely displayed.

## Configuring the agent on Windows systems

You can configure the agent on Windows systems by using the IBM Performance Management window.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, complete these steps:
  - a) Double-click the **Monitoring Agent for MySQL** template.
  - b) In the **Monitoring Agent for MySQL** window, specify an instance name and click **OK**.
3. In the **Monitoring Agent for MySQL** window, complete these steps:
  - a) In the **IP Address** field, enter the IP address of a MySQL server that you want to monitor remotely. If the agent is installed on the server to be monitored, retain the default value.
  - b) In the **JDBC user name** field, enter the name of a MySQL server user. The default value is root.
  - c) In the **JDBC password** field, type the password of a JDBC user.
  - d) In the **Confirm JDBC password** field, type the password again.
  - e) In the **JDBC Jar File** field, click **Browse** and locate the directory that contains the MySQL connector Java file and select it.
  - f) Click **Next**.
  - g) In the **JDBC port number** field, specify the port number of the JDBC server. The default port number is 3306.
  - h) From the **Java trace level** list, select a trace level for Java. The default value is Error.
  - i) Click **OK**.  
The instance is displayed in the **IBM Performance Management** window.
4. Right-click the **Monitoring Agent for MySQL** instance, and click **Start**.

**Remember:** To configure the agent again, complete these steps in the **IBM Performance Management** window:

- a. Stop the agent instance that you want to configure.

- b. Right-click the **Monitoring Agent for MySQL** instance, and click **Reconfigure**.
- c. Repeat steps [3](#) and [4](#).

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring the agent on Linux systems

You run the configuration script to configure the agent on Linux systems.

### Procedure

1. Run the following command:

```
install_dir/bin/mysql-agent.sh config instance_name
```

Where *instance\_name* is the name you want to give to the instance, and *install\_dir* is the installation directory for the MySQL agent.

2. When you are prompted to enter a value for the following parameters, press Enter to accept the default value, or specify a different value and press enter.

- IP Address
- JDBC user name
- JDBC password
- Re-type:JDBC password
- JDBC Jar File
- JDBC port number (Default port number is 3306.)
- Java trace level (Default value is `ERROR`.)

For information about the configuration parameters, see [“Configuring the agent by using the silent response file”](#) on page 407.

3. Run the following command to start the agent.

```
install_dir/bin/mysql-agent.sh start instance_name
```

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring the agent by using the silent response file

Use the silent response file to configure the agent without responding to prompts when you run the configuration script. You can use the silent response file for configuring the agent on both Windows and Linux systems.

### About this task

The silent response file contains the configuration parameters. You edit the parameter values in the response file, and run the configuration script to create an agent instance and update the configuration values.

### Procedure

**Windows** This procedure assumes the following default path where the agent is installed:

**Windows** C:\IBM\APM

**Linux** opt/ibm/apm/agent

If the agent is installed at a different path, substitute the path in the instructions, and edit the **AGENT\_HOME** parameter in the silent response file to specify the path where the agent is installed.

1. In a text editor, open the response file that is available at the following path:

**Linux** `install_dir/samples/mysql_silent_config.txt`

**Windows** `install_dir\samples\mysql_silent_config.txt`

Where `install_dir` is the installation directory of the MySQL agent.

2. In the response file, specify a value for the following parameters:

- For the **Server Name** parameter, specify the IP address of a MySQL server that you want to monitor remotely. Otherwise, retain the default value as `localhost`.
- For the **JDBC user name** parameter, retain the default user name value of `root` or specify the name of a user with privileges to view the `INFORMATION_SCHEMA` tables.
- For the **JDBC password** parameter, enter a JDBC user password.
- For the **JDBC Jar File** parameter, retain the default path if this path to the MySQL connector for the Java jar file is correct. Otherwise, enter the correct path. The connector is available at the following default path:

**Linux** `/usr/share/java/mysql-connector-java.jar`

**Windows** `C:\Program Files (x86)\MySQL\Connector J 5.1.26\mysql-connector-java-5.1.26-bin.jar`

- For the **JDBC port number** parameter, retain the default port number of 3306 or specify a different port number.
  - For the **Java trace level** parameter, retain the default value of `Error` or specify a different level according to the IBM support instructions.
3. Save and close the response file, and run the following command to update the agent configuration settings:

**Linux** `install_dir/bin/mysql-agent.sh config instance_name install_dir/samples/mysql_silent_config.txt`

**Windows** `install_dir\BIN\mysql-agent.bat config instance_name install_dir\samples\mysql_silent_config.txt`

Where `instance_name` is the name that you want to give to the instance, and `install_dir` is the installation directory of MySQL agent.

**Important:** Be sure to include the absolute path to the silent response file. Otherwise, no agent data is displayed in the dashboards.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring NetApp Storage monitoring

You must configure the NetApp Storage agent to monitor the health and performance of NetApp storage systems. You can configure the agent on Windows and Linux systems.

### Before you begin

- Review the hardware and software prerequisites. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\) for the NetApp Storage agent](#).

- Ensure that the following components are installed on your system:
  - OnCommand Unified Manager
  - OnCommand Performance Manager
  - OnCommand API Services

For information about installing these components, see the NetApp documentation.

- Ensure that the versions of the OnCommand API Services, the OnCommand Unified Manager, and the OnCommand Performance Manager are compatible. For example, to configure the OnCommand API Services V1.0, pair the OnCommand Unified Manager V6.2, V6.1, or V6.0 with the OnCommand Performance Manager V1.1. For compatible product versions, see the [Interoperability Matrix Tool](#).
- Ensure that the user, who connects to the OnCommand Unified Manager, has the GlobalRead privilege for the NetApp storage system that is being monitored. Use an existing user ID with this privilege, or create a new user ID. For more information, see the NetApp documentation.
- Ensure that the user, who configures the OnCommand API Services, is an administrator or a monitor. These user types have default permissions to run the rest API.
- Download the NetApp Manageability SDK JAR file (`manageontap.jar`) from the NetApp website and install the file in the monitoring agent `lib` directory by completing the steps that are mentioned in [“Downloading and installing the NetApp Manageability SDK JAR file”](#) on page 409.

#### About this task

The NetApp Storage agent is a multiple instance agent. You must create the first instance, and start the agent manually.

- To configure the agent on Windows systems, you can use the **IBM Performance Management** window or the silent response file.
- To configure the agent on Linux systems, you can run the script and respond to prompts, or use the silent response file.

### Downloading and installing the NetApp Manageability SDK JAR file

The NetApp Storage agent requires the NetApp Manageability SDK JAR file to communicate with the NetApp OCUM server.

#### About this task

After you install the NetApp Storage agent, download the NetApp Manageability SDK JAR file (`manageontap.jar`) from the NetApp website and install the file in the monitoring agent `lib` directory.

#### Procedure

To download and install the NetApp Manageability SDK JAR file, follow these steps:

1. Download the compressed file that contains the JAR file from the following website: <http://communities.netapp.com/docs/DOC-1152>.
2. Extract the compressed file and copy the `manageontap.jar` file to following locations:
  - For 32-bit Windows systems, copy the file to `install_dir/tmaitm6`
  - For 64-bit Windows systems, copy the file to `install_dir/tmaitm6_x64`
  - For 32-bit Linux systems, copy the file to `install_dir/li6263/nu/lib`
  - For 64-bit x86-64 Linux systems, copy the file to `install_dir/lx8266/nu/lib`
  - For 64-bit Linux on System z systems, copy the file to `install_dir/ls3266/nu/lib`

## Configuring the agent on Windows systems

You can configure the agent on Windows operating systems by using the IBM Performance Management window. After you update the configuration values, you must start the agent to save the updated values.

### About this task

The NetApp Storage agent provides default values for some parameters. You can specify different values for these parameters.

### Procedure

To configure the agent on Windows systems, follow these steps:

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Monitoring Agent for NetApp Storage**, and then click **Configure agent**.

**Remember:** After you configure the agent for the first time, the **Configure agent** option is disabled. To configure the agent again, click **Reconfigure**.

3. In the Monitoring Agent for NetApp Storage window, complete the following steps:
  - a) Enter a unique name for the NetApp Storage agent instance, and click **OK**.
  - b) On the **Data Provider** tab, specify values for the configuration parameters, and then click **Next**.
  - c) On the **OnCommand Unified Manager** tab, specify values for the configuration parameters, and then click **Next**.
  - d) On the **OnCommand API Service** tab, specify values for the configuration parameters, and then click **OK**.

For more information about configuration parameters, see the following topics:

- [“Configuration parameters for the data provider” on page 412](#)
  - [“Configuration parameters for the OnCommand Unified Manager” on page 413](#)
  - [“Configuration parameters for the OnCommand API Service” on page 414](#)
4. In the **IBM Performance Management** window, right-click **Monitoring Agent for NetApp Storage**, and then click **Start**.

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

The silent response file contains the agent configuration parameters with default values that are defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters.

After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

To configure the NetApp Storage agent in the silent mode, follow these steps:

1. In a text editor, open the `netapp_storage_silent_config.txt` file that is available at the following path:

- `Linux install_dir/samples/netapp_storage_silent_config.txt`

For example, `/opt/ibm/apm/agent/samples/netapp_storage_silent_config.txt`

- **Windows** `install_dir\samples\netapp_storage_silent_config.txt`

For example, `C:\IBM\APM\samples\netapp_storage_silent_config.txt`

2. In the `netapp_storage_silent_config.txt` file, specify values for all mandatory parameters. You can also modify the default values of other configuration parameters.

For more information, see the following topics:

- [“Configuration parameters for the data provider” on page 412](#)
- [“Configuration parameters for the OnCommand Unified Manager” on page 413](#)
- [“Configuration parameters for the OnCommand API Service” on page 414](#)

3. Save and close the `netapp_storage_silent_config.txt` file, and run the following command:

- **Linux** `install_dir/bin/netapp_storage-agent.sh config instance_name install_dir/samples/netapp_storage_silent_config.txt`

For example, `/opt/ibm/apm/agent/bin/netapp_storage-agent.sh config instance_name /opt/ibm/apm/agent/samples/netapp_storage_silent_config.txt`

- **Windows** `install_dir\bin\netapp_storage-agent.bat config instance_name install_dir\samples\netapp_storage_silent_config.txt`

For example, `C:\IBM\APM\bin\netapp_storage-agent.bat config instance_name C:\IBM\APM\samples\netapp_storage_silent_config.txt`

Where,

**instance\_name**

Name that you want to give to the instance.

**install\_dir**

Path where the agent is installed.

**Important:** Ensure that you include the absolute path to the silent response file. Otherwise, the agent data is not shown in the dashboards.

4. Run the following command to start the agent:

- **Linux** `install_dir/bin/netapp_storage-agent.sh start instance_name`

For example, `/opt/ibm/apm/agent/bin/netapp_storage-agent.sh start instance_name`

- **Windows** `install_dir\bin\netapp_storage-agent.bat start instance_name`

For example, `C:\IBM\APM\bin\netapp_storage-agent.bat start instance_name`

## Configuring the agent on Linux systems

To configure the agent on Linux operating systems, you must run the script and respond to prompts.

### Procedure

To configure the agent on Linux systems, follow these steps:

1. On command line, enter the following command:

```
install_dir/bin/netapp_storage-agent.sh config instance_name
```

For example, `/opt/ibm/apm/agent/bin/netapp_storage-agent.sh config instance_name`

Where,

**instance\_name**

Name that you want to give to the instance.

**install\_dir**

Path where the agent is installed.

2. Respond to the prompts by referring to the following topics:
  - [“Configuration parameters for the data provider” on page 412](#)
  - [“Configuration parameters for the OnCommand Unified Manager” on page 413](#)
  - [“Configuration parameters for the OnCommand API Service” on page 414](#)
3. Run the following command to start the agent:

```
install_dir/bin/netapp_storage-agent.sh start instance_name
```

For example, `/opt/ibm/apm/agent/bin/netapp_storage-agent.sh start instance_name`

## Configuration parameters for the data provider

When you configure the NetApp Storage agent, you can change the default values of the parameters for the data provider. For example, the maximum number of data provider log files, the maximum size of the log file, and the level of detail that is included in the log file.

The following table contains detailed description of the configuration parameters for the data provider.

| <i>Table 52. Name and description of the configuration parameters for the data provider</i> |                                                                                                                                                                                                                                                                        |                        |
|---------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| <b>Parameter name</b>                                                                       | <b>Description</b>                                                                                                                                                                                                                                                     | <b>Mandatory field</b> |
| Instance Name<br>( <b>KNU_INSTANCE_NAME</b> )                                               | The name of the instance.<br><br><b>Restriction:</b> The Instance Name field displays the name of the instance that you specify when you configure the agent for the first time. When you configure the agent again, you cannot change the instance name of the agent. | Yes                    |
| Maximum number of Data Provider log files<br>( <b>KNU_LOG_FILE_MAX_COUNT</b> )              | The maximum number of log files that the data provider creates before it overwrites the previous log files. The default value is 10.                                                                                                                                   | Yes                    |
| Maximum Size in KB of Each Data Provider Log<br>( <b>KNU_LOG_FILE_MAX_SIZE</b> )            | The maximum size in KB that a data provider log file must reach before the data provider creates a new log file. The default value is 5190 KB.                                                                                                                         | Yes                    |



Table 52. Name and description of the configuration parameters for the data provider (continued)

| Parameter name                                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Mandatory field |
|------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Level of Detail in Data Provider Log<br>( <b>KNU_LOG_LEVEL</b> ) | <p>The level of details that you can include in the log file that the data provider creates. The default value is 4. The following values are valid:</p> <ul style="list-style-type: none"> <li>• 1 (Off): No messages are logged.</li> <li>• 2 (Severe): Only errors are logged.</li> <li>• 3 (Warning): All errors and messages that are logged at the Severe level and potential errors that might result in undesirable behavior.</li> <li>• 4 (information): All errors and messages that are logged at the Warning level and high-level informational messages that describe the state of the data provider when it is processed.</li> <li>• 5 (Fine): All errors and messages that are logged at the information level and low-level informational messages that describe the state of the data provider when it is processed.</li> <li>• 6 (Finer): All errors and messages that are logged at the Fine level plus highly detailed informational messages, such as performance profiling information and debug data. Selecting this option can adversely affect the performance of the monitoring agent. This setting is intended only as a tool for problem determination along with IBM support staff.</li> <li>• 7 (Finest): All errors and messages that are logged at the Fine level and the most detailed informational messages that include low-level programming messages and data. Choosing this option might adversely affect the performance of the monitoring agent. This setting is intended only as a tool for problem determination along with IBM support staff.</li> <li>• 8 (All): All errors and messages are logged.</li> </ul> | Yes             |

### Configuration parameters for the OnCommand Unified Manager

When you configure the NetApp Storage agent, you can change the default values of the parameters for the OnCommand Unified Manager (OCUM), such as the IP address of the OCUM server, user name, and password.

The following table contains detailed description of configuration parameters for the data source.

Table 53. Name and description of the configuration parameters for the OnCommand Unified Manager

| Parameter name                                   | Description                                                                                                   | Mandatory field |
|--------------------------------------------------|---------------------------------------------------------------------------------------------------------------|-----------------|
| Server<br>( <b>KNU_DATASOURCE_HOST_ADDRESS</b> ) | The host name or IP address of the NetApp OCUM server that you want to monitor.                               | Yes             |
| User ( <b>KNU_DATASOURCE_USERNAME</b> )          | A user name on the NetApp OCUM server with sufficient privileges to collect data. The default value is admin. | Yes             |
| Password<br>( <b>KNU_DATASOURCE_PASSWORD</b> )   | The password of the user that you specify in the <b>User</b> parameter.                                       | Yes             |

Table 53. Name and description of the configuration parameters for the OnCommand Unified Manager (continued)

| Parameter name                     | Description                                                                                               | Mandatory field |
|------------------------------------|-----------------------------------------------------------------------------------------------------------|-----------------|
| Confirm Password                   | The same password that you specified in the <b>Enter Password</b> parameter.                              | Yes             |
| Protocol (KNU_DATASOURCE_PROTOCOL) | The protocol that you want to use to communicate with the NetApp OCUM server. The default value is HTTPS. | Yes             |

## Configuration parameters for the OnCommand API Service

When you configure the NetApp Storage agent, you can change the default values of configuration parameters for the OnCommand API Service, such as the host address, user name, and password.

The following table contains detailed description of configuration parameters for the data source.

Table 54. Name and description of configuration parameters for the OnCommand API Service

| Parameter name                               | Description                                                                                                 | Mandatory field |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------|-----------------|
| Host Address (KNU_API_SERVICES_HOST_ADDRESS) | The host name or IP address of the OnCommand API service.                                                   | Yes             |
| User (KNU_API_SERVICES_USERNAME)             | A user name with sufficient privileges to connect to the OnCommand API service. The default value is admin. | Yes             |
| Password (KNU_API_SERVICES_PASSWORD)         | The password of the user that you specify in the <b>User</b> parameter.                                     |                 |
| Confirm Password                             | The same password that you specified in the <b>Enter Password</b> parameter.                                | Yes             |

## Configuring Oracle Database monitoring

The Monitoring Agent for Oracle Database provides monitoring capabilities for the availability, performance, and resource usage of the Oracle database. You can configure more than one Oracle Database agent instance to monitor different Oracle databases. Remote monitoring capability is also provided by this agent.

### Before you begin

- Before you configure the Oracle Database agent, you must grant privileges to the Oracle user account that is used by the Oracle Database agent. For more information about privileges, see [Granting privileges to the Oracle Database agent user](#).
- If you are monitoring an Oracle database remotely, the agent must be installed on a computer with either the Oracle database software or the Oracle Instant Client installed.

### About this task

The directions here are for the most current release of the agent, except as indicated. For information about how to check the version of an agent in your environment, see [Agent version](#).

For general Oracle database performance monitoring, the Oracle Database agent provides monitoring for the availability, performance, resource usage, and activities of the Oracle database, for example:

- Availability of instances in the monitored Oracle database.
- Resource information such as memory, caches, segments, resource limitation, tablespace, undo (rollback), system metric, and system statistics.
- Activity information, such as OS statistics, sessions, contention, and alert log.

The Oracle Database agent is a multiple-instance agent. You must create the first instance and start the agent manually. Additionally, each agent instance can monitor multiple databases.

The Managed System Name for the Oracle Database agent includes a database connection name that you specify, an agent instance name that you specify, and the host name of the computer where the agent is installed. For example, `pc:connection_name-instance_name-host_name:SUB`, where *pc* is your two character product code and *SUB* is the database type (Possible values are RDB, ASM, or DG). The Managed System Name is limited to 32 characters. The instance name that you specify is limited to 23 characters, minus the length of your host name and database connection. For example, if you specify **dbconn** as your database connection name, **Oracle02** as your agent instance name, and your host name is *Prod204a*, your managed system name is `RZ:dbconn-oracle02-Prod204a:RDB`. This example uses 22 of the 23 characters available for the database connection name, agent instance name, and host name.

- If you specify a long instance name, the Managed System name is truncated and the agent code does not display correctly.
- The length of the *connection\_name*, *instance\_name*, and *hostname\_name* variables are truncated when they exceed 23 characters.
- To avoid a subnode name that is truncated, change the subnode naming convention by setting the following environment variables: **KRZ\_SUBNODE\_INCLUDING\_AGENTNAME**, **KRZ\_SUBNODE\_INCLUDING\_HOSTNAME**, and **KRZ\_MAX\_SUBNODE\_ID\_LENGTH**.
- If you set **KRZ\_SUBNODE\_INCLUDING\_AGENTNAME** to NO, the subnode ID part of the subnode name does not include the agent instance name. For example,
  - Default subnode name: *DBConnection-Instance-Hostname*
  - Subnode name with environment variable set to NO: *DBConnection-Hostname*
- If you set **KRZ\_SUBNODE\_INCLUDING\_HOSTNAME** to NO, the subnode ID part of the subnode name does not include the host name. For example,
  - Default subnode name: *DBConnection-Instance-Hostname*
  - Subnode name with environment variable set to NO: *DBConnection-Instance*

## Procedure

1. To configure the agent on Windows systems, you can use the **IBM Performance Management** window or the silent response file.
  - [“Configuring the agent on Windows systems” on page 416.](#)
  - [“Configuring the agent by using the silent response file” on page 423.](#)
2. To configure the agent on Linux and UNIX systems, you can run the script and respond to prompts, or use the silent response file.
  - [“Configuring the agent by responding to prompts” on page 419.](#)
  - [“Configuring the agent by using the silent response file” on page 423.](#)

## What to do next

For advanced configuration only, the Oracle database administrator must enable the Oracle user to run the `krzgrant.sql` script to access the database, see [Running the krzgrant.sql script](#).

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

If you are unable to view the data in the agent dashboards, first check the server connection logs and then the data provider logs. The default paths to these logs are as follows:

- **Linux** | **UNIX** /opt/ibm/apm/agent/logs
- **Windows** C:\IBM\APM\TMAITM6\_x64\logs

## Configuring the agent on Windows systems

You can configure the agent on Windows operating systems by using the **IBM Performance Management** window. After you update the configuration values, start the agent to apply the updated values.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Cloud App Management**.
2. In the **IBM Performance Management** window, right-click the **Monitoring Agent for Oracle Database** template, and then click **Configure agent**.

**Remember:** After you configure an agent instance for the first time, the **Configure agent** option is disabled. To configure the agent instance again, right-click on it and then click **Reconfigure....**

3. In the Monitoring Agent for Oracle Database window, complete the following steps:
  - a) Enter a unique instance name for the Monitoring Agent for Oracle Database instance, and click **OK**.
4. On the Default Database Configuration pane of the **Configure ITCAM Extended Agent for Oracle Database** window, perform the following steps:
  - a) Enter the **Default Username**. This is the default database user ID for database connections.  
This user ID is the ID that the agent uses to access the monitored database instance. This user ID must have select privileges on the dynamic performance views and tables that are required by the agent.
  - b) Enter the **Default Password**. This is the password that is associated with the specified default database user ID.
  - c) Enter the **Oracle JDBC Jar File**. This is the full path to the Oracle JDBC driver jar file used to communicate with the Oracle database.  
The Oracle Java Database Connectivity (JDBC) driver that supports the Oracle database versions monitored by the Oracle agent must be available on the agent computer.
  - d) If you need to set advanced configuration options, check **Show advanced options** otherwise, proceed to [step 5](#).
  - e) Net Configuration Files Directories can be left blank and the default directory is used. Only one directory is supported.  
This setting contains the Oracle database net configuration file or files. The directory is defined by the *TNS\_ADMIN* environment variable for each Oracle database instance. The default directory is %ORACLE\_HOME%\NETWORK\ADMIN. If this item is not configured, the default directory is used. To disable the use of the default directory, set the following agent environment variable to false: KRZ\_LOAD\_ORACLE\_NET=false.
  - f) Leave the Customized SQL definition file name blank. It is not used.
  - g) Choose whether the default dynamic listener is configured at this workstation.  
The default dynamic listener is (PROTOCOL=TCP) (HOST=localhost) (PORT=1521). If the default dynamic listener is configured at this workstation, set this value to Yes.
  - h) Click **Next**.
5. On the **Instance configuration** pane of the **Configure ITCAM Extended Agent for Oracle Database** window, perform the following steps:

This is where the actual database connection instances are defined. You need to add at least one. This is also where you edit and delete database connection instances. If multiple database connection

instance configurations exist, use the **Database connections** option to choose the instance to edit or delete.

a) Press **New** in the **Database connections** section.

b) Enter a **Database Connection Name** as an alias for the connection to the database.

This alias can be anything that you choose to represent the database connection with the following restrictions. Only letters from the Latin alphabet (a-z, A-Z), Arabic numerals (0-9), the underline character (  ), and the hyphen-minus character (-) can be used in the connection name. The maximum length of a connection name is 25 characters.

c) Choose a **Connection Type**

1) (Optional) Basic

The default and most common connection type is **Basic**. If you are unsure which connection type you need, it is suggested that you choose this connection type.

a) Select the **Basic** connection type when the target monitored database is a single instance, such as a standard file system instance or an ASM single instance.

b) Enter the **Hostname** as the host name or IP address for the database.

c) Enter the **Port** number that is used by the database.

d) Select either **Service Name** or **SID**.

i. When **Service Name** is selected, enter the name of the service that is a logical representation of a database, a string that is the global database service name.

A service name is a logical representation of a database, which is the way that a database is presented to clients. A database can be presented as multiple services and a service can be implemented as multiple database instances. The service name is a string that is the global database name, that is, a name composed of the database name and domain name, entered during installation or database creation. If you are not sure what the global database name is, then you can obtain it from the value of the `SERVICE_NAMES` parameter in the initialization parameter file.

ii. When **SID** is selected, enter the Oracle System Identifier that identifies a specific instance of a running database.

This is the Oracle System Identifier that identifies a specific instance of a database.

Proceed to [step 5d](#).

2) (Optional) TNS

a) Select the **TNS** connection type if the `ORACLE_HOME` system environment variable is set and the TNS alias for the target monitored database is defined in the `$ORACLE_HOME/network/admin/tnsnames.ora` file.

b) Enter the **TNS alias** name.

Proceed to [step 5d](#).

3) (Optional) Advanced

a) Select the **Advanced** connection type when there is more than one Oracle Instance across multiple physical nodes for the target monitored database. For example, an ASM with Real Applications Cluster (RAC) database.

b) Enter the **Oracle Connection String**.

This attribute supports all Oracle Net naming methods as follows:

- SQL Connect URL string of the form://host:port/service name. For example, //d1sun242:1521/bjava21.
- Oracle Net keyword-value pair. For example,

```
(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp) (HOST=d1sun242) (PORT=1521))
(CONNECT_DATA=(SERVICE_NAME=bjava21)))
```

- **TNSNAMES** entries, such as **inst1**, with the **TNS\_ADMIN** or **ORACLE\_HOME** environment variable set and the configuration files configured.

Proceed to [step 5d](#).

- d) Check **Use a different user name and password** for this connection to use different credentials than the default credentials that you set in [step 4a](#) and [step 4b](#). Otherwise, proceed to [step 5g](#).
- e) Enter the **Database Username** for this connection.  
This user ID is the ID that the agent uses to access the monitored database instance. This user ID must have select privileges on the dynamic performance views and tables that are required by the agent.
- f) Enter the **Database Password**. The password that is associated with the specified database user ID.
- g) Select a **Role** that matches the permissions that are granted to the database connection's credentials.  
The role is the set of privileges to be associated with the connection. For a user that was granted the SYSDBA system privilege, specify a role that includes that privilege. For ASM instances, use the **SYSDBA** or **SYSASM** role.
- h) Check **Show remote log monitoring options** if you monitor remote Oracle alert logs from this agent instance, otherwise proceed to [step 5k](#).
- i) Enter a path or use **Browse** to select the **Oracle Alert log file paths**.

The absolute file paths of mapped alert log files for remote database instances in this database connection. The agent monitors alert logs by reading these files. Usually found at \$ORACLE\_BASE/diag/rdbms/DB\_NAME/SID/trace/alert\_SID.log. For example, if the DB\_NAME and SID are both db11g and ORACLE\_BASE is /home/dbowner/app/oracle, then the alert log would be found at /home/dbowner/app/oracle/diag/rdbms/db11g/db11g/trace/alert\_db11g.log.

**Windows** If the Oracle Database agent runs and reads the alert log files through the network, the remote file path must follow the universal naming convention for Windows systems. For example, \\tivx015\path>alert\_orcl.log.

**Windows**

**Important:** Enter the path and alert log file name together. A mapped network driver is not supported for the alert log path.

**Linux** **UNIX** If the Oracle Database agent is on a remote server, a locally mounted file system is required to monitor its remote alert logs.

**Windows** Multiple files are separated by a semicolon (;).

**Linux** **UNIX** Multiple files are separated by a colon (:).

Each file is matched to a database instance by using the alert\_instance.log file name pattern or if it is unmatched, it is ignored.

Local database instance alert log files are discovered automatically.

- j) Select or enter the **Oracle Alert Log File Charset**. This is the code page of the mapped alert log files.

If this parameter is blank, the system's current locale setting is used, for example:

- ISO8859\_1, ISO 8859-1 Western European encoding
- UTF-8, UTF-8 encoding of Unicode
- GB18030, Simplified Chinese GB18030 encoding

- CP950, Traditional Chinese encoding
- EUC\_JP, Japanese encoding
- EUC\_KR, Korean encoding

For the full list of all the supported code pages, see the [ICU supported code pages](#).

- k) Click **Apply** to save this database connection instance's settings in the **Database connections** section.
  - l) (Optional) Test the new database connection.
    - 1) Select the new database connection in the **Database connections** section.
    - 2) Click **Test connection**.
    - 3) Observe the results in the **Test connection** result window.
      - Example successful **Test Result**:
 

```
Testing connection config1 ...
Success
```
      - Example unsuccessful **Test Result**:
 

```
Testing connection config1 ...
KBB_RAS1_LOG; Set MAXFILES to 1
ORA-12514: TNS:listener does not currently know of service requested in connect
descriptor
Failed
```
  - m) Click **Next**.
6. Read the information on the **Summary** pane of the **Configure ITCAM Extended Agent for Oracle Database** window, then click **OK** to finish configuration of the agent instance.
  7. In the **IBM Performance Management** window, right-click **Monitoring Agent for Oracle Database**, and then click **Start**.

### What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent by responding to prompts

To configure the agent on Linux and UNIX operating systems, run the command line configuration script and respond to its prompts.

### Procedure

1. Open the `install_dir/bin` directory, where `install_dir` is the installation directory for the Oracle Database agent.
2. (Optional) To list the names of any existing configured agent instances, run the following command: `./cinfo -o rz`.
3. To configure the Oracle Database agent, run the following command: `./oracle_database-agent.sh config instance_name`.
4. When prompted to Edit 'Monitoring Agent for Oracle Database' settings, press **Enter**. The default value is Yes.
5. To enter the Default Database Configuration information, perform the following steps:

**Note:** The Default Database Configuration section is not the database connection instance configuration. It is a template section for setting what is used as the default values when you add the actual database connection instance configurations, which begin in [step 6](#).

- a) When prompted for the Default Username, type the default database user ID for database connections and press **Enter**.

This user ID is the ID that the agent uses to access the monitored database instance. This user ID must have select privileges on the dynamic performance views and tables that are required by the agent.

- b) When prompted to Enter Default Password, type the password that is associated with the specified default database user ID, and press **Enter**. Then, if prompted, confirm the password.
- c) Enter the **Oracle JDBC Jar File**. This is the full path to the Oracle JDBC driver jar file used to communicate with the Oracle database.

The Oracle Java Database Connectivity (JDBC) driver that supports the Oracle database versions monitored by the Oracle agent must be available on the agent computer.

- d) Net Configuration Files Directories can be left blank and the default directory is used. If the Oracle agent version is 6.3.1.10, you can enter multiple net configuration file directories by using `Windows ";"` or `Linux | UNIX ":"` to separate the directories. For Oracle agent version 8.0, only one directory is supported. Press **Enter**.

This setting contains the Oracle database net configuration file or files. The directory is defined by the `TNS_ADMIN` environment variable for each Oracle database instance. The default directory is

`Linux | UNIX $ORACLE_HOME/network/admin` or `Windows %ORACLE_HOME%\NETWORK\ADMIN`. If this item is not configured, the default directory is used. To disable the use of the default directory, set the following agent environment variable to false:  
`KRZ_LOAD_ORACLE_NET=false`.

- e) Choose whether the default dynamic listener is configured at this workstation, and press **Enter**.

The default dynamic listener is `(PROTOCOL=TCP) (HOST=localhost) (PORT=1521)`. If the default dynamic listener is configured at this workstation, set this value to True.

- f) Leave the Customized SQL definition file name blank. It is not used.

6. You are prompted to Edit 'Database Connection' settings after seeing the following output on the screen:

```
Instance Configuration :  
Summary :  
Database Connection :
```

**Note:** This step is where the actual database connection instances are defined. You need to add at least one. This is also where you edit and delete database connection instances. If multiple database connection instance configurations exist, use the Next option to skip the instances that do not need to be edited or deleted until you arrive at the instance you need to edit or delete.

7. To add a new database connection, type 1, and press **Enter**.
8. To enter the database connection information, perform the following steps:

- a) When prompted for the Database Connection Name, type an alias for the connection to the database and press **Enter**.

This alias can be anything that you choose to represent the database connection with the following restrictions. Only letters, Arabic numerals, the underline character, and the minus character can be used in the connection name. The maximum length of a connection name is 25 characters.

- b) When prompted for the Connection Type, select one of the following types of connection:

- 1) (Optional) Basic

The default and most common connection type is **Basic**. If you are unsure which connection type you need, it is suggested that you choose this connection type.

- a) Select the **Basic** connection type if the target monitored database is a single instance, such as a standard file system instance or an ASM single instance.
- b) When prompted for the Hostname, type the host name or IP address for the Oracle database, and press **Enter**.
- c) When prompted for the Port, type the port number, and press **Enter**.
- d) Enter one of the next two settings. Either Service Name or SID.



- i. (Optional) When prompted for the `Service Name`, type the name of the service that is a logical representation of a database, a string that is the global database service name, press **Enter** and proceed to [step 8c](#).

A service name is a logical representation of a database, which is the way that a database is presented to clients. A database can be presented as multiple services and a service can be implemented as multiple database instances. The service name is a string that is the global database name, that is, a name composed of the database name and domain name, entered during installation or database creation. If you are not sure what the global database name is, then you can obtain it from the value of the `SERVICE_NAMES` parameter in the initialization parameter file. This parameter can be left blank if you set the `SID` in step “[8.b.i.4.b](#)” on page 421.

- ii. (Optional) When prompted for the `SID`, type the Oracle System Identifier that identifies a specific instance of a running database, press **Enter** and proceed to [step 8c](#).

This parameter is the Oracle System Identifier that identifies a specific instance of a database. If `Service Name` was defined in step “[8.b.i.4.a](#)” on page 421, you can leave this item blank.

## 2) (Optional) TNS

- a) Select the **TNS** connection type when the `ORACLE_HOME` system environment variable is set and the TNS alias for the target monitored database is defined in the `$ORACLE_HOME/network/admin/tnsnames.ora` file.
- b) Type the TNS alias name, and press **Enter** and proceed to [step 8c](#).

## 3) (Optional) Advanced

- a) Select the **Advanced** connection type when there is more than one Oracle Instance across multiple physical nodes for the target monitored database. For example, an ASM with Real Applications Cluster (RAC) database.
- b) Type the Oracle connection string, press **Enter** and proceed to [step 8c](#).

This attribute supports all Oracle Net naming methods as follows:

- SQL Connect URL string of the form: `//host:port/service name`. For example, `//d1sun242:1521/bjava21`.
- Oracle Net keyword-value pair. For example,

```
(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp) (HOST=d1sun242) (PORT=1521))
(CONNECT_DATA=(SERVICE_NAME=bjava21)))
```

- **TNSNAMES** entries, such as **inst1**, with the `TNS_ADMIN` or `ORACLE_HOME` environment variable set and the configuration files configured.

**Note:** The description that is shown during command-line configuration might have a backslash before colons (\:) and before equal sign symbols (\=). Do not type backslashes in the connection string. They are displayed in the description to escape the normal behavior of interpreting the equals sign as part of a command, and instead interpret it merely as text.

- c) Proceed to [step 8c](#).
- c) When prompted for the `Database Username`, type the database user ID for the connection, and press **Enter**.  
For standard file system instances, this user ID must have select privileges on the dynamic performance views and tables that are required by the agent.  
For ASM instances, use an account with the **SYSDBA** or **SYSASM** role. For example, the `sys` account.
- d) When prompted to `Enter Database Password`, type the password that is associated with the specified database user ID.
- e) When prompted for `Role`, choose the role that matches the permissions that are granted to the specified user ID, and press **Enter**.

The role is the set of privileges to be associated with the connection. For a user that was granted the SYSDBA system privilege, specify a role that includes that privilege.

For ASM instances, use the **SYSDBA** or **SYSASM** role.

- f) When prompted for Oracle Alert Log File Paths (including alert log file name), type the alert log paths, and press **Enter**.

This parameter is for any absolute file paths of mapped alert log files for remote database instances in this database connection. The agent monitors alert logs by reading these files. Usually found at \$ORACLE\_BASE/diag/rdbms/DB\_NAME/SID/trace/alert\_SID.log. For example, if the DB\_NAME and SID are both db11g and ORACLE\_BASE is /home/dbowner/app/oracle, then the alert log would be found at /home/dbowner/app/oracle/diag/rdbms/db11g/db11g/trace/alert\_db11g.log.

**Windows** If the Oracle Database agent runs and reads the alert log files across the network, the remote file path must follow the universal naming convention for Windows systems. For example, \\tivx015\path\alert\_orcl.log.

**Important:** Enter the path and alert log file name together. A mapped network driver is not supported for the alert log path.

**Linux** **UNIX** If the Oracle Database agent runs, a locally mounted file system is required for remote alert logs.

**Windows** Multiple files are separated by a semicolon (;).

**Linux** **UNIX** Multiple files are separated by a colon (:).

Each file is matched to a database instance by using the alert\_instance.log file name pattern or if it is unmatched, it is ignored.

Local database instance alert log files can be discovered automatically.

- g) When prompted for the **Oracle Alert Log File Charset**, type the code page of the mapped alert log files, and press **Enter**.

If this parameter is blank, the system's current locale setting is used, for example:

- ISO8859\_1, ISO 8859-1 Western European encoding
- UTF-8, UTF-8 encoding of Unicode
- GB18030, Simplified Chinese GB18030 encoding
- CP950, Traditional Chinese encoding
- EUC\_JP, Japanese encoding
- EUC\_KR, Korean encoding

For the full list of all the supported code pages, see the [ICU supported code pages](#).

9. When prompted again to Edit 'Database Connection' settings, you see the name of the database connection that you set in [step 8a](#). You can edit it again or delete it. If you have more than one database connection instance that is already configured, use **Next** to step through them.
10. (Optional) To add another database connection to monitor multiple database instances with this agent instance, type 1, press **Enter**, and return to [Step 8](#).
11. When you are finished modifying database connections, type 5, and press **Enter** to exit the configuration process.
12. To start the agent, enter:  
`install_dir/bin/oracle_database-agent.sh start instance_name.`

### What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance or update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

The silent response file contains the agent configuration parameters with default values that are defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters.

After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

1. Open the `oracle_silent_config.txt` file in a text editor:

- **Linux** | **UNIX** `install_dir/samples/oracle_database_silent_config.txt`.
- **Windows** `install_dir\samples\oracle_database_silent_config.txt`

2. For **Default Username**, type the name of the default database user for database connections that are created for this agent instance. For example, `KRZ_CONN_USERID=user1`.

**Note:** This user must have sufficient privileges to complete the tasks that this agent performs while it is connected to the database, such as querying tables.

3. For **Default Password**, you must enter the password that is associated with the specified default database user. For example, `KRZ_CONN_PASSWORD=Password`.

4. Enter the **Oracle JDBC Jar File**. This is the full path to the Oracle JDBC driver jar file used to communicate with the Oracle database.

The Oracle Java Database Connectivity (JDBC) driver that supports the Oracle database versions monitored by the Oracle agent must be available on the agent computer.

5. Net Configuration Files Directories can be left blank and the default directory is used. The Oracle Database agent uses this file path to obtain the `tnsnames.ora` file. This directory is defined by the `TNS_ADMIN` environment variable for each Oracle database instance. The default directory is **Linux** | **UNIX** `$ORACLE_HOME/network/admin` or **Windows** `%ORACLE_HOME%\NETWORK\ADMIN`. If you enter this setting with multiple net configuration file directories, use **Windows** `","` or **Linux** | **UNIX** `":"` to separate the directories.

If you are monitoring Oracle databases remotely, you can copy net configuration files from the remote system to the system where the agent is installed. Also, you can merge the content of net configuration files on the remote system to the net configuration files on the system where the agent is installed.

6. For **Dynamic listener**, check if the default dynamic listener is configured. The default dynamic listener is `(PROTOCOL=TCP)(HOST=localhost)(PORT=1521)`. If the default dynamic listener is configured, set this value to `TRUE` as shown here; `KRZ_DYNAMIC_LISTENER=TRUE`.

The valid values are `TRUE` and `FALSE`.

7. Leave the Customized SQL definition file name blank. It is not used.

8. Beginning here the actual database connection instances are defined. You need to add at least one. Entries for one instance are given in the `oracle_silent_config.txt` with the instance name `config1`. If you change the instance name, be sure to change all references.

This alias can be anything that you choose to represent the database connection with the following restrictions. Only letters, Arabic numerals, the underline character, and the minus character can be used in the connection name. The maximum length of a connection name is 25 characters.

9. For **Connection Type**, specify one of the following connection types: **Basic**, **TNS**, or **Advanced**. For example, `KRZ_CONN_TYPE.config1=Basic`.

10. For the connection type that you selected in the previous step, specify the required parameters:

### Basic

- For **Hostname**, specify the host name or the IP address of the Oracle database, for example: **KRZ\_CONN\_HOST.config1=** hostname.
- For **Port**, specify the Listener port for the Oracle database, for example: **#KRZ\_CONN\_PORT.config1=** 1521.
- For **Service Name**, specify the logical representation of the database by using a string for the global database name, for example: **KRZ\_CONN\_SERVICE.config1=** orcl.

**Important:** If you do not define the Service Name, you must specify the Oracle System Identifier (SID).

For the **Oracle System Identifier (SID)**, specify an SID that identifies a specific instance of a running database, for example: **KRZ\_CONN\_SID.config1=** sid.

### TNS

For **TNS alias**, specify the Network alias name from the `tnsnames.ora` file. For example, **KRZ\_CONN\_TNS.config1=** tnsalias.

### Advanced

For **Oracle Connection String**, specify the database connection string for OCI. For example, **KRZ\_CONN\_STR.config1=** //host:port/service

This string supports all Oracle Net naming methods as shown here.

- For an SQL Connect URL string:

```
//host:[port][/]service name]
```

- For an Oracle Net keyword-value pair:

```
"(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP) (HOST=d1sun242) (PORT=1521))
(CONNECT_DATA=(SERVICE_NAME=bjava21)))"
```

This string also supports **TNSNAMES** entries, for example, **inst1** where the `TNS_ADMIN` or the `ORACLE_HOME` environment variable is set and the configuration files are configured.

**Important:** This attribute applies only to the advanced type of connection.

11. For **Database Username**, you can specify the name of the database user for the connection, for example: **KRZ\_CONN\_USERID=**UserID.

This user must have sufficient privileges to complete the tasks that the agent requires while it is connected to the database, for example, creating, editing, and deleting tables.

If this field is empty, the agent uses the default user name in the default database configuration section. If **Database Username** was not configured, the default user name is used for this connection.

12. For **Database Password**, you can specify the password that is associated with the specified database user, for example: **KRZ\_CONN\_PASSWORD=**Password.

If this field is empty, the agent uses the default password in the default database configuration section. If **Database Password** was not configured, the default password is used for this connection.

13. For **Role**, you can specify the set of privileges that are associated with the connection, for example: **KRZ\_CONN\_MODE.config1=**DEFAULT.

The valid values include `SYSDBA`, `SYSOPER`, `SYSASM`, and `DEFAULT`.

For a user that is granted the `SYSDBA` system privilege, you can specify a connection that includes this privilege. If this item is not defined, you can assign the `DEFAULT` role to the user.

14. For **Oracle Alert Log File Paths**, when the alert log file name is included, you can specify the absolute file path of the mapped alert log files for the remote database instances in this database connection. For example, **KRZ\_LOG\_PATHS.config1=**AlertLogPath.

**Windows** Use a semicolon (;) to separate the multiple files.

**Linux** **UNIX** Use a colon (:) to separate the multiple files.

Each file is matched to a database instance by the `alert_instance.log` file name pattern. Alternatively, it is ignored if it is not matched.

The local database instance alert log files are discovered automatically.

If **Oracle Alert Log File Paths** was not configured, the Alert Log is not available.

15. For **Oracle Alert Log File Charset**, you can specify the code page of the mapped alert log files. For example, `KRZ_LOG_CHARSET.config1= CharSet`

If this field is empty, the system's current locale setting is used as shown here:

```
ISO8859_1: ISO 8859-1 Western European encoding
UTF-8: UTF-8 encoding of Unicode
GB18030: Simplified Chinese GB18030 encoding
CP950: Traditional Chinese encoding
EUC_JP: Japanese encoding
```

16. Save and close the `oracle_database_silent_config.txt` file. Then, enter:  
`install_dir/bin/oracle_database-agent.sh config instance_name install_dir/samples/oracle_database_silent_config.txt`  
where `instance_name` is the name that you want to give to the instance.
17. To start the agent, enter:  
`install_dir/bin/oracle_database-agent.sh start instance_name.`

### What to do next

Log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176.](#)

## Granting privileges to the Oracle Database agent user

After you install the agent, you must grant privileges to the Oracle user account that is used by the Oracle Database agent.

You can grant privileges for the following users:

- Standard file system (non-ASM) instance users
- ASM with RAC instance non-SYS users

### Granting privileges to users for standard file system instances

For standard file system instances, the Oracle user ID that the Oracle Database agent uses must have select privileges on the dynamic performance views, tables, and data dictionary views that are required by the agent. It must also have other Oracle object and system privileges that are necessary to run some database commands.

### Procedure

1. (Optional) If an Oracle database user ID does not exist, create this ID by using Oracle facilities and running the following command: `create user UserName identified by Password`
2. Grant select privileges for the dynamic performance views, tables, and data dictionary views to the Oracle user ID that you created by running the `krzgrant.sql` script that is provided with the Oracle Database agent. This step must be done before you configure the agent. For directions about how to customize and run the `krzgrant.sql` script, see [“Customizing the krzgrant.sql script” on page 426](#) and [“Running the krzgrant.sql script” on page 426.](#)

**Note:** The select privileges for the dynamic performance views, tables, and data dictionary views rely on the capabilities of the Oracle database in specific application environments. You can grant authorized Oracle privileges to the Oracle database user ID only for the dynamic performance views, tables, and data dictionary views that are used by the Oracle Database agent.

3. Grant other Oracle object privileges and system privileges to the Oracle user ID that the Oracle Database agent uses by using Oracle facilities.

### Customizing the `krzgrant.sql` script

If you do not want to allow Oracle authorized select privileges on some dynamic performance views, tables, and data dictionary views in the `krzgrant.sql` script, you can customize the `krzgrant.sql` script before running it.

**Note:** The agent instance checks all default privileges in the `krzgrant.sql` script and reports an agent event with a lack of privileges when the agent starts. You can disable privilege checking by using the following variable setting: `KRZ_CHECK_ORACLE_PRIVILEGE=FALSE`. The test connection step of GUI configuration checks all Oracle privileges that are defined in the `krzgrant.sql` file. If you confirm that the Oracle user has the correct privileges, ignore that checking privileges fails in the test connection step.

Edit the `krzgrant.sql` file in a plain text editor to remove or add the '--' prefix at the beginning of grant statements to skip the granting execution for those unauthorized Oracle tables or views.

For example, change the following lines:

```
execute immediate 'grant select on DBA_HIST_SNAPSHOT to '||userName;
execute immediate 'grant select on DBA_HIST_SQLSTAT to '||userName;
execute immediate 'grant select on DBA_HIST_SQLTEXT to '||userName;
execute immediate 'grant select on DBA_HIST_SQL_PLAN to '||userName;
execute immediate 'grant select on DBA_HIST_SYSMETRIC_SUMMARY to '||userName;
```

to these lines:

```
-- execute immediate 'grant select on DBA_HIST_SNAPSHOT to '||userName;
-- execute immediate 'grant select on DBA_HIST_SQLSTAT to '||userName;
-- execute immediate 'grant select on DBA_HIST_SQLTEXT to '||userName;
-- execute immediate 'grant select on DBA_HIST_SQL_PLAN to '||userName;
-- execute immediate 'grant select on DBA_HIST_SYSMETRIC_SUMMARY to '||userName;
```

### Granting privileges to non-SYS users for ASM instances

You must connect to ASM instances that are using the SYSDBA and SYSASM roles for users. If you do not want to use the SYS account to connect to ASM instances, create a user account and grant the SYSDBA and SYSASM roles to the account.

#### Procedure

1. Run the following commands to create a user account and grant roles:

- Log in to the ASM database with the SYSASM role to create a new user for an agent and grant the SYSDBA role or SYSASM role:

a. `create user UserName identified by Password`

b. `grant sysdba to UserName`

or

`grant sysasm to UserName`

2. When you create the ASM connection in the configuration window, specify the *UserName* user and the SYSDBA or SYSASM role.

### Running the `krzgrant.sql` script

#### Before you begin

- If you do not run the `krzgrant.sql` script, an event is raised in the agent event workspace.

After the installation, you can find the `krzgrant.sql` script in the following directory:

- **Windows** `install_dir\TMAITM6_X64`

- **Linux** | **UNIX** `install_dir/architecture/rz/bin`

where:

**install\_dir**

Installation directory for the Oracle Database agent.

**architecture**

The IBM Cloud App Management or Cloud App Management system architecture identifier. For example, lx8266 represents Linux Intel v2.6 (64-bit). For a complete list of the architecture codes, see the `install_dir/registry/archdsc.tbl` file.

The **krzgrant.sql** script has the following usage: `krzgrant.sql user_ID temporary_directory`

where:

**user\_ID**

The ID of the Oracle user. This user ID must be created before you run this SQL file. Example value: `tivoli`.

**temporary\_directory**

The name of the temporary directory that contains the `krzagent.log` output file of the **krzgrant.sql** script. This directory must exist before you run this SQL script. Example value: `install_dir/tmp`.

You must have the Oracle database administrator (DBA) authorization role and write permission to the temporary directory to perform the following procedure.

**Procedure**

1. From the command line, run the commands to set environment variables.

- **Windows**

```
SET ORACLE_SID= sid
SET ORACLE_HOME= home
```

- **Linux** | **UNIX**

```
ORACLE_SID = sid
export ORACLE_SID
ORACLE_HOME = home
export ORACLE_HOME
```

where:

**sid**

Oracle system identifier, which is case-sensitive.

**home**

Home directory for the monitored Oracle instance.

2. From the same command-line window where you set environment variables, start the Oracle SQL Plus or an alternative tool that you use to issue SQL statements.
3. Log on to the Oracle database as a user that has Oracle DBA privileges.
4. Go to the directory that contains the **krzgrant.sql** script and run the following command to grant select privileges:

```
@krzgrant.sql user_ID temporary_directory
```

The output is logged in the `krzagent.log` file in the temporary directory. This log records the views and tables to which the Oracle Database agent is granted select privileges.

After the privileges are successfully granted, you can configure and start the Oracle Database agent.

## Configuring OS monitoring

The Monitoring Agent for Linux OS, Monitoring Agent for UNIX OS, and Monitoring Agent for Windows OS agents are configured automatically. However, you can configure log file monitoring for the OS agents so that you can monitor application log files. Also, you can run the OS agents as a non-root user.

### Run OS agents as a non-root user

You can run the Monitoring Agent for Windows OS, Monitoring Agent for UNIX OS, and Monitoring Agent for Linux OS as a non-root user.

To run the Windows OS agent as a non-root user, see [“Run the Monitoring Agent for Windows OS as a non-root user” on page 428](#).

To run the Monitoring Agent for UNIX OS and Monitoring Agent for Linux OS agents as a non-root user, see [“Starting agents as a non-root user” on page 230](#).

#### Restriction:

When you run the OS agent as a non-root user, the agent cannot access `/proc/pid/status`, and therefore cannot report the following attributes:

- -User CPU Time (UNIXPS.USERTIME)
- -System CPU Time (UNIXPS.SYSTEMTIM)
- -Total CPU Time (UNIXPS.TOTALTIME)
- -Thread Count (UNIXPS.THREADCNT)
- -Child User CPU Time (UNIXPS.CHILDUTIME)
- -Child System CPU Time (UNIXPS.CHILDSTIME)
- -Total Child CPU Time (UNIXPS.CHILDTIME)
- -Wait CPU Time (UNIXPS.WAITCPUTIM)
- -Terminal (UNIXPS.USERTTY)

These attributes are not visible in the Cloud APM console but are available to create thresholds.

### Run the Monitoring Agent for Windows OS as a non-root user

You can run the Windows OS agent as a non-root user. However, some functions are unavailable.

When you run the Windows OS agent as a non-root user, some functions are unavailable in the following attribute groups, if they are owned solely by the administrator account:

- Registry
- File Trend
- File Change

Remote deployment of other agents is not available because administrator rights are required to install the new agents.

For Agent Management Services, the watchdog cannot stop or start any agent that does not have privileges to stop or start.

To create a non-root user, create a new Limited (non-root) user and set up registry permissions for the new user as in the following example:

- Full access to `HKEY_LOCAL_MACHINE\SOFTWARE\Candle`
- Read access to `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib`

The user that starts the Monitoring Agent for Windows OS – Primary service must have rights to manage the Monitoring Agent for Windows OS - Watchdog service. The user that starts the Monitoring Agent for Windows OS - Watchdog service must also have rights to manage any services that are managed by the Agent Management Services, including the Monitoring Agent for Windows OS – Primary service. To grant



users the authority to manage system services in Windows, use security templates, group policy, or edit the Subinacl.exe file. For more information, see the following Microsoft documentation: <http://support.microsoft.com/kb/325349> (<http://support.microsoft.com/kb/325349>).

The following example shows how to grant users the authority to manage system services by using security templates:

1. Click **Start > Run**, enter `mmc` in the Open box, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**.
3. Click **Add > Security Configuration and Analysis**, and then click **Add** again.
4. Click **Close** and then click **OK**.
5. In the console tree, right-click **Security Configuration and Analysis**, and then click **Open Database**.
6. Specify a name and location for the database, and then click **Open**.
7. In the **Import Template** dialog box that is displayed, click the security template that you want to import, and then click **Open**.
8. In the console tree, right-click **Security Configuration and Analysis**, and then click **Analyze Computer Now**.
9. In the **Perform Analysis** dialog box that is displayed, accept the default path for the log file that is displayed in the Error log file path box. Otherwise, specify the location that you want. Click **OK**.
10. After the analysis is complete, configure the service permissions as follows:
  - a. In the console tree, click **System Services**.
  - b. In the right pane, double-click the Monitoring Agent for Windows OS - Primary service.
  - c. Select the **Define this policy in the database** check box, and then click **Edit Security**.
  - d. To configure permissions for a new user or group, click **Add**.
  - e. In the **Select Users, Computers, or Groups** dialog box, type the name of the user or group that you want to set permissions for, and then click **OK**. In the **Permissions for User or Group** list, select the **Allow** check box (next to **Start**). Stop and pause permission is selected by default, so that the user or group can start, stop, or pause the service.
  - f. Click **OK** twice.
11. Repeat step 10 to configure the service permissions for the Monitoring Agent for Windows OS - Watchdog service.
12. To apply the new security settings to the local computer, right-click **Security Configuration and Analysis**, and then click **Configure Computer Now**.

**Note:** You can use also the Secedit command to configure and analyze system security. For more information about Secedit, click **Start > Run**, enter `cmd`, and then click **OK**. At the command prompt, type `secedit /?`, and then press **ENTER**. When you use this method to apply settings, all the settings in the template are reapplied. This method might override other previously configured file, registry, or service permissions.

The following example shows how to set the Monitoring Agent for Windows OS and Watchdog services to log on as a non-root user by using the Windows Services console:

1. Click **Start > Run**, enter `services.msc`, and then click **OK**.
2. Select **Monitoring Agent for Windows OS - Primary**.
3. Right-click **Properties**.
4. Verify the startup type as being Automatic.
5. Select the **Log On** tab, and then select **Log on as "This account"** and supply the ID and password. Click **OK**.
6. Select **Monitoring Agent for Windows OS - Watchdog**.
7. Right-click **Properties**.
8. Verify the startup type as being Manual.

9. Select the **Log On** tab, and then select **Log on as "This account"** and supply the ID and password. Click **OK**.

## Configure OS agent log file monitoring

The Monitoring Agent for Linux OS, Monitoring Agent for UNIX OS, and Monitoring Agent for Windows OS agents are configured automatically. However, you can configure log file monitoring for the OS agents so that you can monitor application log files.

After the agents filter the log data, the data is sent in the form of a log event to the Cloud App Management console.

### Adding or removing log file monitoring configuration for the OS agents

You add log file monitoring configuration for the OS agents so the OS agents can filter log file data. Additionally, you can also remove the log file monitoring configuration for the OS agents, if necessary.

#### Before you begin

The OS agents now include a sample `regex1.conf` file and a `regex1.fmt` file that you can view before you configure `.conf` and `.fmt` files. The files are located here:

- On UNIX/LINUX: `<install_dir>/samples/logfile-monitoring`
- On windows: `<install_dir>\samples\logfile-monitoring`

Use a text editor to create a configuration `.conf` file and a format `.fmt` file. For more information about the content of these files, see [“Configuration file”](#) on page 432 and [“Format file ”](#) on page 441.

#### About this task

To enable the OS agents to monitor log files, you must create and save the configuration and format files in the directory: `<install_dir>/localconfig/<pc>/log_discovery`.

#### Procedure

Adding the log file monitoring configuration for the OS agents

1. Add log file monitoring configuration and format file for the OS agents to the following location:  
`install_dir/localconfig/pc/log_discovery`  
where `install_dir` might be `/opt/ibm/apm/agent` on UNIX/Linux or `C:\IBM\APM` on Windows.  
and `pc` is `lz`, `ux`, or `nt`.

Removing the log file monitoring configuration for the OS agents

2. To stop log file monitoring, remove the configuration and format file from `install_dir/localconfig/pc/log_discovery`.

#### Important:

After you remove the log monitoring configuration, the log file monitoring resource remains and it stays online until you restart the OS agent.

#### Viewing log file monitoring content

You can view the log file monitoring configuration for the OS agents that you deployed to monitor log files.

#### Procedure

1. Drill down to the OS agent resources (Linux Systems, Unix Systems, Windows Systems) in the OS agent dashboard.
2. Expand the **OS agent monitor logs** widgets.

The screenshot shows the 'OS agent monitor logs' interface. It is divided into two main sections: 'File status' and 'Profile events'. Both sections have a search filter box.

**File status table:**

| Filename                                           | Table name | Processed records | Matched recor... |
|----------------------------------------------------|------------|-------------------|------------------|
| <input checked="" type="radio"/> /var/log/messages | MonLnxLog  | 1759321           | 1759321          |
| <input type="radio"/> /var/log/secure              | MonLnxLog  | 17                | 17               |

**Profile events table:**

| Timestamp     | Messages                                                                                | Occurrence |
|---------------|-----------------------------------------------------------------------------------------|------------|
| 10/3, 6:38 AM | Scope libcontainer-31819-systemd-test-default-dependencies.scope has no PIDs. Refusing. | 1484       |
| 10/3, 6:38 AM | Scope libcontainer-5943-systemd-test-default-dependencies.scope has no PIDs. Refusing.  | 1          |
| 10/3, 6:39 AM | Scope libcontainer-5943-systemd-test-default-                                           | 1502       |

## Displaying log file monitoring events

After you configure the OS agent to monitor your application log files, you can create thresholds to raise alarms on the log file conditions that you want to be alerted of.

### Procedure

1. Refer to the Managing thresholds section for creating threshold.
2. Select metrics that begins with **Log File Profile**, **Log File Status**, **Log Profile Events**.

### Results

When the specified condition becomes true, the log file event that triggers the alert is displayed in the Events tab.

### Log file monitoring environment variables

You can set environment variables for log file monitoring in the OS agent environment files.

Set the following environment variables and replace *KPC* with the OS agent code where *PC* is the two character agent code, for example, *klz* is the code for the Linux OS agent.

#### **KPC\_FCP\_LOG**

This variable is available in the `install_dir/config/.pc.environment` file. The default value is `True` and you use it to enable or disable the log monitoring feature.

#### **KPC\_FCP\_LOG\_PROCESS\_MAX\_CPU\_PCT**

This setting is the maximum allowable percentage of all system CPU that the agent uses over a 1-minute interval. Valid values are 5 - 100. The default value is 100. This setting is associated with the CPU throttling feature. If you specify a value less than 5, the minimum value of 5 is used.

#### **KPC\_FCP\_LOG\_PROCESS\_PRIORITY\_CLASS**

This setting is the operating system scheduler priority for the process. A is lowest, C is the operating system default, and F is the highest priority. The setting is one of the following values: A, B, C, D, E, F. These values are superseded by any values that you specify in the `.conf` file.

#### **KPC\_FCP\_LOG\_SEND\_EVENTS**

The default setting is `True` and it is used by the OS agent to send events to the Cloud APM server.

#### **KPC\_FCP\_LOG\_SEND{EIF\_EVENTS**

The default setting is `True`. If this option is set to `Yes` the agent sends event data to the Cloud APM server or to any EIF receiver such as the OMNIBus EIF probe. If the option is set to `No`, the agent does not send the event data. The setting of this option is global and applies to all monitoring profiles.

**Note:** The EIF receiver consumes events, otherwise problems might occur when the agent cache fills.

#### **KPC\_FCP\_TRUNCATE\_HOSTNAME\_IN\_LONG\_SUBNODE\_MSN**

OS agents with log file event monitoring have a subnode limitation. To manage log file events, the subnode MSN has the following structure: `UX:CITRAHOSTNAME_PROFILENAME`. The maximum size limitation for the subnode name is 32 characters. If the built subnode MSN name is too long and it is more than 32 characters, it is truncated to 32 characters. This name corresponds to the substring that is taken from the Profile Name.

In the OS agent configuration file, use the following variables to manage the profile names that are too long:

- UNIX OS agent: `KUX_FCP_TRUNCATE_HOSTNAME_IN_LONG_SUBNODE_MSN=true`
- Linux OS agent: `KLZ_FCP_TRUNCATE_HOSTNAME_IN_LONG_SUBNODE_MSN=true`
- Windows OS agent: `KNT_FCP_TRUNCATE_HOSTNAME_IN_LONG_SUBNODE_MSN=true`

For example, if you have an agent that is called `aixhost_nc123456789A`, which is 20 characters in length, `CTIRAHOSTNAME=aixhost_nc123456789A` is 20 characters.

and you have two profiles that are called:

```
ProfileLong12A (14 characters)
ProfileLong12B (14 characters)
```

the following related subnode MSNs are expected:

```
UX:aixhost_nc123456789A_ProfileLong12A (38 characters)
UX:aixhost_nc123456789A_ProfileLong12B (38 characters)
```

However, the subnode MSNs are truncated to the 32 character limitation so the resulting names are the same for both:

```
UX:aixhost_nc123456789A_ProfileL
UX:aixhost_nc123456789A_ProfileL
```

To truncate `CTIRAHOSTNAME` instead of the Profile Name, set the `Kpc_FCP_TRUNCATE_HOSTNAME_IN_LONG_SUBNODE_MSN=true` variable.

For example, if  $n$  is the length of the Profile Name, such as 14, the substring for the MSN name that relates to `CTIRAHOSTNAME` is truncated to  $32-n-3$  characters, so the `CTIRAHOSTNAME` variable is: `aixhost_nc1234`. Then, the distinguished subnode MSNs are:

```
UX:aixhost_nc1234_ProfileLong12A
UX:aixhost_nc1234_ProfileLong12B
```

### Configuration file

OS agents use a configuration file `t` that is read by the agent when it starts. The file contains configuration options and filters. You must create this configuration file and configure the agent instance to use it.

The configuration file is monitored for changes to its time stamp every 60 seconds thereafter. If the time stamp of the file changes, the agent reinitializes its configuration dynamically, without requiring a restart. For more information, see [“Changing the agent configuration and format files” on page 444](#).

The `.conf` file for the OS agent accepts these options:

#### codepage

This parameter is the code page of the monitored file. Use this parameter in the configuration file when the code page of the monitored file is different from the code page of the system. Specify the code page of the monitored file, for example, `ibm-5348_P100-1997`, `UTF-16`, or `UTF-8`.

#### ConfigFilesAreUTF8=Y

This parameter specifies that the configuration file and format file are in UTF-8. Use this parameter if the encoding of the configuration files is UTF-8 and the system code page is not. The default is that the agent assumes the system encoding.

#### DupDetectionKeyAttributes

A comma-separated list of Cloud APM attributes that is used to determine which events are duplicates. If all the named attributes are the same in two events, then those two events are considered duplicates. This option applies only to events. For more information, see [Chapter 24, “Event filtering and summarization,” on page 1425](#).

#### Note:

1. The attribute names are case-sensitive, so you must enter the names exactly as described.
2. If you do not provide a list of attributes, the values default to `Class` and `Logname`.

## **ENFORCE\_STRICT\_TEC\_COMPATIBILITY**

This parameter refers to all white space characters in the log data to ensure that the characters are respected. For example, when you use a format such as "%s %s" to extract information from log messages, the OS agent matches not only a literal space but also any other white space characters that are present such as tabs and carriage returns.

When this parameter is not set, the default behavior of the OS agent when it matches a Tivoli Enterprise Console® style format string is to match as much of the input text as it can, while it processes the format from left to right.

For example, for the %s : %s format string and the one : two : three input string, the OS agent default assigns one . two to the first parameter (corresponding to the first %s) and it assigns three to the second parameter.

### **Note:**

1. This parameter does not apply to format statements that use the regular expression syntax.
2. Setting this parameter has a performance impact. To give greater control over the behavior and performance of matching, avoid setting this parameter and use regular expressions instead.

## **EventSummaryInterval**

Specifies the number of seconds during which the agent searches for duplicate events to suppress. Set this parameter to a positive integer. This option applies only to events. For more information, see [Chapter 24, "Event filtering and summarization," on page 1425](#).

## **EventFloodThreshold**

Specifies which events are sent when duplicate events are detected. Set this parameter to `send_none`, `send_all`, `send_first`, or a positive integer. This option applies only to events. For more information, see [Chapter 24, "Event filtering and summarization," on page 1425](#).

## **EventMaxSize**

Specifies in bytes, the maximum size of a generated event. If specified, this parameter is used in two places:

1. The parameter can be used by the agent to set the size of a buffer that is used to process events. If not set, this buffer defaults to a size of 16384 bytes. If the buffer is set too small, events are truncated and can be discarded.
2. The parameter can be used by the EIF sender to set the size of a buffer that is used to send events to an EIF receiver, such as the OMNIbus EIF probe. If not set, this buffer defaults to a size of 4096 bytes. If the buffer is set too small, events are discarded.

## **FileComparisonMode**

Specifies which log files are monitored when more than one file matches a wildcard pattern. The following values are available:

### **CompareByAllMatches**

This value is the default behavior. All files that match the wildcard pattern that is specified in `LogSources` are monitored.

### **CompareByLastUpdate**

Of the files that match the wildcard pattern that is specified in `LogSources`, the file with the most recently updated time stamp is monitored.

### **CompareBySize**

Of the two or more files that match the file name pattern criteria, the larger file is selected for monitoring. Do not use `CompareBySize` with multiple matching files that are being updated at the same time and increasing their file sizes. If the largest file is subject to frequent change, monitoring might continually restart at the beginning of the newly selected file. Instead, use `CompareBySize` for a set of matching files where only one is active and being updated at any specific time.

### **CompareByCreationTime**

Of the files that match the wildcard pattern that is specified in `LogSources`, the file with the most recently created time stamp is monitored. This value has the following restrictions:

- The value is applicable only to Windows operating systems because UNIX and Linux operating systems do not store a true creation time for files.
- The value is not supported for remote files that you monitor by using the Secure Shell (SSH) File Transfer Protocol.

**Tip:** The `CompareByLastUpdate`, `CompareBySize`, and `CompareByCreationTime` values can all be used for rolling log files. `CompareByLastUpdate` is typically used for these files.

### **FQDomain**

Specifies how and if the agent sets a domain name:

- If set to `yes`, the agent determines the system domain name.
- If set to `no`, the agent does not set a domain name. The `fqhostname` attribute is assigned a blank string.
- If set so that it does not contain a `yes` or `no` value, the domain name is accepted as the value and it is appended to the host name.

For more information, see [“Format file”](#) on page 441.

### **IncludeEIFEventAttr**

The agent includes a large attribute that is called *EIFEvent*, which is a representation of the event that is sent through the Event Integration Facility if that feature is enabled. The information that is contained in the *EIFEvent* attribute can also be found in other attributes. Its large size made it problematic, thus it was disabled by default. Setting this value to `y`, reenables the *EIFEvent* attribute.

**Note:** Using this attribute might cause thresholds to fail if you have large events. A large event in this context is an event where the total number of bytes that is required to contain all values for all attributes and their names results in a string longer than 3600 bytes.

### **LognameIsBaseName**

When set to `y`, the value of the `Logname` attribute is the base name of the log file in which the event was found. This option applies only to Performance Management events. The path is removed. For example, `/data/logs/mylog.log` becomes `mylog.log`. If this value is set to `n`, then you get the full path. However, because the attribute is limited to 64 characters, setting it to `n` means that the name is truncated if it is longer. For this reason, the default value is `y`. To see the full path name in a longer attribute, you can specify it in the mappings section of a format in the `.fmt` file, for example, `filename FILENAME CustomSlot1`. The mapping completes the slot that is named `filename` with the full path of the file in which the event was found and maps it into `CustomSlot1`, which is 256 characters.

### **LogSources**

Specifies the text log files to poll for messages. The complete path to each file must be specified, and file names must be separated by commas. Within each file name, you can also use an asterisk (\*) to represent any sequence of characters, or a question mark (?) to represent any single character. For example, `mylog*` results in polling all log files whose names begin with `mylog`, whereas `mylog???` results in polling all log files whose names consist of `mylog` followed by exactly 3 characters. These wildcard characters are supported only within the file name; the path must be explicitly specified.

If you want to use regular expressions or pattern matching in the path, see the [RegexLogSources](#) description.

A log file source is not required to exist when the agent is started; the log file is polled when it is created.

### **NewFilePollInterval**

Specifies the frequency, in seconds, that the agent checks for new files to monitor. For example, if a file name specified by the *LogSources* or *RegexLogSources* configuration file settings does not yet exist when the agent starts, it checks again for the existence of the files after this interval.

### **NumEventsToCatchUp**

Specifies the event in the log that the agent starts with. This option provides some flexibility if the source that is being monitored is new or the agent is stopped for an extended time. The following values are valid:

**Note:** For text files, values 0 and -1 apply. For Windows Event Log, values 0, -1, and n apply.

**0**

Start with the next event in the logs. This value is the default.

**-1**

When set to -1, the agent saves its place in the file that is being monitored. It saves its place so that when the agent is stopped and later restarted, it can process any events that are written to the log while it was stopped. The agent otherwise ignores events that arrived while it was stopped and restarts from the end of the file. This setting does not apply to pipes, or syslog monitoring on UNIX and Linux systems.

**n**

Set to a positive integer. Starts with the *n*th event from the most current event in the logs; that is, start *n* events back from the most current event in the logs. If *n* is greater than the number of events that are available, all the events that are available are processed.

**Note:** You can use the *n* value only for Windows Event Log. The *n* value is ignored when UseNewEventLogAPI is set to *y*.

### **PollInterval**

Specifies the frequency, in seconds, to poll each log file that is listed in the LogSources option for new messages. The default value is 5 seconds.

If you upgraded a Windows Event Log adapter from a previous release and you have a value that is set for PollingInterval in the Windows registry, you must specify the PollInterval option in the agent configuration file with the same value that is used in the Windows registry. This rule applies only if you are replacing a Tivoli Enterprise Console OS agent that had values in the registry.

### **ProcessPriorityClass**

Specifies the process priority for the agent. You can adjust this value to improve system performance if the agent processes large volumes of events and is using too many processor resources. The possible values are:

- A - Very low priority
- B - Low priority
- C - Typical priority
- D - Above typical priority
- E - High priority
- F - Very high priority
- USE\_CONF\_FILE\_VALUE - Use the value that is specified in the configuration file. This value is the default.

### **RegexLogSources**

Specifies the text log files to poll for messages. It differs from the LogSources option in that regular expression meta characters can be used in the base name portion of the file name and in one subdirectory of the file name. This difference provides greater flexibility than the LogSources option in describing multiple files to monitor in multiple directories.

For example, specifying `/var/log/mylog*` for the LogSources statement is identical to using the dot (`.`) meta character followed by an asterisk (`*`) meta character to form `/var/log/mylog.*` in the RegexLogSources statement. This type of qualifier results in polling all log files in the `/var/log` directory whose base names begin with `mylog` and are followed by zero or more characters.

A `/var/log/mylog.+` qualifier results in polling all log files in the `/var/log` directory whose names begin with `mylog` and are followed by one or more characters.

Similar to LogSources, the complete path to each file must be specified and the file names must be separated by commas. However, the comma is also a valid character inside a regular expression. To distinguish between a comma that is used as part of a regular expression and one that is used to separate file names, commas that are used as part of a regular expression must be escaped with the backslash (`\`) character.

For example, if you want to search for logs that match either of the following regular expressions, `/logs/*.log` and `/other/logs/[a-z]{0,3}.log`, you must escape the comma in the `{0,3}` clause of the second expression so the agent does not mistake it for the beginning of a new expression: `RegexLogSources=/logs/*.log,/other/logs/[a-z]{0\,3}.log`

If meta characters are used in the path name, the meta characters can be used in only one subdirectory of the path. For example, you can specify `/var/log/[0-9\.]*/mylog.*` to have meta characters in one subdirectory. The `[0-9\.]*` results in matching any subdirectory of `/var/log` that consists solely of numbers and dots (`.`). The `mylog.*` results in matching any file names in those `/var/log` subdirectories that begin with `mylog` and are followed by zero or more characters.

Because some operating systems use the backslash (`\`) as a directory separator it can be confused with a regular expression escape meta character. Because of this confusion, forward slashes must always be used to indicate directories. For example, Windows files that are specified as `C:\temp\mylog.*` might mean the `\t` is a shorthand tab character. Therefore, always use forward slashes (`/`) for all operating systems directory separators. For example, `C:/temp/mylog.*` represents all files in the `C:/temp` directory that start with `mylog`.

If more than one subdirectory contains meta characters, a trace message is also issued. For example, `c:[0-9\.]*/temp.files/mylog.*` has two subdirectories with meta characters. `[0-9\.]*` is the first subdirectory with meta characters and `temp.files` is the second subdirectory that used a dot (`.`) meta character. In this case, the agent assumes that the first subdirectory with the meta character is used and the subsequent directories with meta characters are ignored.

#### **SubnodeName**

A string value that can be used to override the default name that is assigned to a monitoring profile subnode. By default the subnode name that is assigned to a monitoring profile corresponds to the base name of the configuration file that is used for that profile. By using this setting, a different subnode name can be assigned.

#### **SubnodeDescription**

A string value that can be used to assign a value to the *Subnode Description* attribute of *LFAProfiles*.

#### **UnmatchLog**

Specifies a file to log discarded events that cannot be parsed into an event class by the agent. The discarded events can then be analyzed to determine whether modifications to the agent format file are required. Events that match a pattern that uses `*DISCARD*` do not appear in the unmatch log because they did match a pattern.

This option is used in a test environment to validate the filters in the format file. This option fills up your file system if you leave it on for extended periods.

### **Options for remote log file monitoring by using SSH**

Other than **SshHostList**, which is a list, all options can have only one value, which is applied to all remote hosts that are specified in **SshHostList**.

Only text log files are supported. AIX error report, syslog, and Windows Event Log are not supported.

**Tip:** You can set up syslog to write its output to a text log file and then remotely monitor that text file with the OS agent.

#### **SshAuthType**

Must be set to either *PASSWORD* or *PUBLICKEY*. If set to *PASSWORD*, the value of **SshPassword** is treated as the password to be used for SSH authentication with all remote systems. If set to *PUBLICKEY*, the value of **SshPassword** is treated as the pass phrase that controls access to the private key file. If set to *PUBLICKEY*, **SshPrivKeyfile** and **SshPubKeyfile** must also be specified.

#### **SshHostList**

A comma-separated list of remote hosts to monitor. All log files that are specified in the **LogSources** or **RegexLogSources** statements are monitored on each host that is listed here. If *localhost* is one of the specified host names, the agent monitors the same set of files directly on the local system. When you specify *localhost*, SSH is not used to access the files on the local system; the log files are read directly.



## SshPassword

When the value of **SshAuthType** is *PASSWORD*, this value is the account password of the user that is specified in **SshUserid**. You can supply the account password in clear text, or you can supply a password that is encrypted with the IBM Tivoli Monitoring CLI **itmpwdsnmp** command. For more information about how to encrypt a password by using the **itmpwdsnmp** command, see [“Remote log file monitoring: Encrypting a password or pass phrase”](#) on page 449.

When the value of **SshAuthType** is *PUBLICKEY*, this value is the pass phrase that decrypts the private key that is specified by the **SshPrivKeyfile** parameter. You can supply the pass phrase in clear text, or you can supply a pass phrase that is encrypted with the IBM Tivoli Monitoring CLI **itmpwdsnmp** command. For more information about how to encrypt a password by using the **itmpwdsnmp** command, see [“Remote log file monitoring: Encrypting a password or pass phrase”](#) on page 449.

**Note:** If the value of **SshAuthType** is *PUBLICKEY*, and you configured SSH not to require a pass phrase, **SshPassword** must be set to null. To set **SshPassword** to null, the entry in the configuration file is:

```
SshPassword=
```

## SshPort

A TCP port to connect to for SSH. If not set, defaults to 22.

## SshPrivKeyfile

If **SshAuthType** is set to *PUBLICKEY*, this value must be the full path to the file that contains the private key of the user that is specified in **SshUserid**, and **SshPubKeyfile** must also be set. If **SshAuthType** is not set to *PUBLICKEY*, this value is not required and is ignored.

## SshPubKeyfile

If **SshAuthType** is set to *PUBLICKEY*, this value must be the full path to the file that contains the public key of the user that is specified in **SshUserid**, and **SshPrivKeyfile** must also be set. If **SshAuthType** is not set to *PUBLICKEY*, this value is not required and is ignored.

## SshUserid

The user name on the remote systems, which the agent uses for SSH authentication.

## Option that is supported on UNIX and Linux systems only

Linux | UNIX

### AutoInitSyslog

If this option is set to Yes, the agent automatically configures the syslog facility to write a standard set of events to a pipe that the agent monitors. By enabling this setting, you can monitor syslog events without maintaining and rolling over log files. If this option is not set in the configuration file, it is the same as being set to No.

**Restriction:** This option is not supported for remote log file monitoring.

## Options that are supported on Windows systems only

Windows

### NTEventLogMaxReadBytes

If you are using the older NT Event Log interface (**UseNewEventLogAPI** is not set to y) to read event log data on a Windows system, the agent reads up to this number of bytes each time it checks the event log for new data. Setting the value to 0 causes the agent to attempt to read all new data, as it did in earlier releases. This activity can occupy the agent for a considerable amount of time on a system with many events. The default value is 655360. When set, the agent might not stop at exactly the value that is specified, but rather at the nearest multiple of an internal buffer size to this value.

### PreFilter

Specifies how events in a Windows Event Log are filtered before agent processing. **PreFilter** statements are used by **PreFilterMode** when the filters determine which events are sent from an event log to the agent. An event matches a **PreFilter** statement when each *attribute=value*

specification in the `PreFilter` statement matches an event in the event log. A `PreFilter` statement must contain at least the log specification and can contain up to three more specifications, which are all optional: event ID, event type, and event source. The order of the attributes in the statement does not matter.

The `PreFilter` statement has the following basic format:

```
PreFilter:Log=log_name;EventId=value; EventType=value;Source=value;
```

You can specify multiple values for each attribute by separating each value with a comma.

Each `PreFilter` statement must be on a single line.

`PreFilter` is not mandatory. All Windows log events are sent to the agent if prefilters are not specified and `PreFilterMode=OUT`.

### PreFilterMode

This option applies only to Windows Event Log. The option specifies whether Windows systems log events that match a `PreFilter` statement are sent (`PreFilterMode=IN`) or ignored (`PreFilterMode=OUT`). Valid values are `IN`, `in`, `OUT`, or `out`. The default value is `OUT`.

`PreFilterMode` is optional; if `PreFilterMode` is not specified, only events that do not match any `PreFilter` statements are sent to the agent.

**Note:** If you set `PreFilterMode=IN`, you must also define the `PreFilter` statements.

### SpaceReplacement

Set to `TRUE` by default for Windows Event Log (Windows Server 2008 only) but not for previous versions of Event Log. When `SpaceReplacement` is `TRUE`, any spaces in the security ID, subsource, Level, and keywords fields of the event log messages are replaced with underscores (`_`). When `SpaceReplacement` is `FALSE`, any spaces in the security ID, subsource, Level, and keywords fields of the event log messages remain unchanged. For more information about this option, see [Chapter 25, “Windows Event Log,”](#) on page 1427.

### UseNewEventLogAPI

When set to `y` on Windows systems, uses the new Windows Event Log interface for event logs. The option is supported only on Windows 2008 and later. The option is needed to access many of the new event logs that debuted in Windows 2008 and the applications that run on it. The option is ignored on earlier versions of Windows and on UNIX and Linux. For more information about this option, see [Chapter 25, “Windows Event Log,”](#) on page 1427.

### WINEVENTLOGS

Controls which Windows event logs are monitored.

The `WINEVENTLOGS` statement is a comma-delimited list with no spaces. For more information, see [Chapter 25, “Windows Event Log,”](#) on page 1427.

**Note:** Any carriage returns, tabs, or new lines in Windows events are replaced by spaces.

### Option that is supported on AIX systems only

AIX

#### AIXErrptCmd

An **errpt** (error report) command string that the agent runs can be supplied here. The command output is fed into the stream of log data that is being monitored.

For example, the following command causes the agent to search for the `mmddhhmmyy` string and replace it with the actual date and time on startup. Only the first occurrence of the string is replaced.

```
AIXErrptCmd=errpt -c -smddhhmmyy
```

Although you can supply your own `errpt` command, you must use the `-c` (concurrent mode) option so that the command runs continuously. You cannot use the `-t` option or the following options that result in detailed output: `-a`, `-A`, or `-g`.

The data stream is the standard output from the `errpt` command, so regular expressions in the `.fmt` file must be written to match. For example, the data output might be:

```
IDENTIFIER  TIMESTAMP  T C RESOURCE_NAME  DESCRIPTION
F7FA22C9   0723182911 I O SYSJ2         UNABLE TO ALLOCATE SPACE IN FILE SYSTEM
2B4F5CAB   1006152710 U U ffdc         UNDETERMINED ERROR
2B4F5CAB   1006152610 U U ffdc         UNDETERMINED ERROR
```

A sample format that picks up the data rows, but not the header, is:

```
REGEX GenericErrpt
^([A-F0-9]{8})+([0-9]{10}) ([A-Z]) ([A-Z]) (\S+) +(.*)$
Identifier $1 CustomSlot1
Timestamp $2 CustomSlot2
T          $3 CustomSlot3
C          $4 CustomSlot4
Resource   $5 CustomSlot5
msg        $6
END
```

For more information, see *Monitoring an AIX Binary Log* in the [IBM Agent Builder User's Guide](#).

### Options that apply only when events are being forwarded to EIF

**Important:** These options apply to EIF events sent directly to Operations Analytics - Log Analysis, OMNIbus, or any other generic EIF receiver. The options are not intended for use with the Cloud APM server.

#### BufferEvents

Specifies how event buffering is enabled. The possible values are:

- **YES** - Stores events in the file that is specified by the `BufEvtPath` option (This value is the default).
- **MEMORY\_ONLY** - Buffers events in memory.
- **NO** - Does not store or buffer events.

#### BufEvtPath

Specifies the full path name of the agent cache file. If this path is not rectified the default is:

- **AIX** `/etc/Tivoli/tec/cache`
- **Windows** `\etc\Tivoli\tec\cache`

**Note:** If events are being forwarded to more than one server, a `BufEvtPath` value must be specified for each forwarding channel. An index number is appended to the `BufEvtPath` name for each additional entry. For example, use `BufEvtPath1` to indicate the path name of the agent cache file for forwarding to the first extra server. The value that is set in each `BufEvtPath` must be unique.

#### BufEvtMaxSize

Specifies the maximum size, in KB, of the agent cache file. The default value is 64. The cache file stores events on disk when the `BufferEvents` option is set to Yes. The minimum size for the file is 8 KB. File sizes specified less than this level are ignored, and 8 KB is used. The value that you specify for the maximum file size does not have an upper limit.

**Note:** If the cache file exists, you must delete the file for option changes to take effect.

#### NO\_UTF8\_CONVERSION

Specifies whether the Event Integration Facility encodes event data in UTF-8. When this option is set to YES, the EIF does not encode event data in UTF-8. The data is assumed to already be in UTF-8 encoding when passed to the EIF. However, a prefix is added to the flag to indicate that the data is in UTF-8 encoding (if the flag does not exist at the beginning of the event data). The default value is NO.

#### MaxEventQueueDepth

This value indicates the maximum number of events that can be queued for forwarding. When the limit is reached, each new event that is placed in the queue bumps the oldest event from the queue. If not specified, the default value is 1000. This setting applies to all forwarding channels if `NumAdditionalServers` is used.

### **NumAdditionalServers**

This entry is required if you want to forward events to more than one Netcool/OMNIBus ObjectServer. Its value is used to indicate the number of servers that events are forwarded to. Valid values are 1 - 8.

### **ServerLocation**

Specifies the name of the host on which the event server is installed. Specify host name or IP address. Use the dotted format for IP address. You can specify failover values such as `ServerLocation1=2.3.4.5,2.3.4.6` for the server locations if you want to. If you specify failover values for *ServerLocation*, you must also specify an extra *ServerPort* value for each *ServerLocation*.

**Note:** If events are being forwarded to more than one server, a *ServerLocation* value must be specified for each server. An index number is appended to the *ServerLocation* name for each additional entry. For example, use *ServerLocation1* to specify the name of the host on which the first extra server is installed.

### **ServerPort**

Specifies the port number on which the EIF receiver listens for events. The *ServerPort* option can contain up to eight values, which are separated by commas. If failover values are specified for *ServerLocation*, you must set an equivalent *ServerPort* value. The *ServerPort* is not used when the *TransportList* option is specified.

**Note:** If events are being forwarded to more than one server, a *ServerPort* value must be specified for each server. An index number is appended to the *ServerPort* name for each additional entry. For example, use *ServerPort1* to specify the port number on which the EIF receiver listens for events for the first extra server.

### **TransportList**

Specifies the user-supplied names of the transport mechanisms, which are separated by commas. When a transport mechanism fails for sender applications, the API uses the following transport mechanisms in the order that is specified in the list. For receiving applications, the API creates and uses all the transport mechanisms. The transport type and channel for each *type\_name* must be specified by using the *Type* and *Channels* keywords:

#### ***type\_name*Type**

Specifies the transport type for the transport mechanism that is specified by the *TransportList* option. SOCKET is the only supported transport type.

The server and port for each *channel\_name* are specified by the *ServerLocation* and *ServerPort* options.

#### ***type\_name*Channels**

##### ***channel\_name*Port**

Specifies the port number on which the transport mechanisms server listens for the specified channel (set by the *Channel* option). When this keyword is set to zero, the portmapper is used. This keyword is required.

##### ***channel\_name*PortMapper**

Enables the portmapper for the specified channel.

##### ***channel\_name*PortMapperName**

Specifies the name of the portmapper if the portmapper is enabled.

##### ***channel\_name*PortMapperNumber**

Specifies the ID that is registered by the remote procedure call.

##### ***channel\_name*PortMapperVersion**

Specifies the version of the portmapper if the portmapper is enabled.

##### ***channel\_name*ServerLocation**

Specifies the name of the event server and the region where the server for transport mechanisms is located for the specified channel. The channel is set by the *Channel* option. This keyword is required.

The configuration file accepts generic EIF options when used directly with OMNIBus. These options operate only over an EIF connection to OMNIBus. They do not affect events that are sent to the Cloud APM server. For more information about these EIF options, see [EIF keywords](#).

### Format file

OS agents extract information from system log messages and then match different log messages to event classes. A format file serves as a lookup file for matching log messages to event classes, which tells the event class what to read, what to match, and how to format the data.

When the format file is used as a lookup file, all format specifications in the file are compared from the beginning to the end of the file. When two classes match or when a message has multiple matching classes, the first expression from the end that matches is used. If no match is found, the event is discarded. A discarded event is written to the unmatched log if it is defined in the `.conf` file.

The regular expression syntax that you use to create patterns to match log messages and events is described. Regular expression-filtering support is provided by using the International Components for Unicode (ICU) libraries to check whether an attribute value that is examined matches the specified pattern.

For more information about using regular expressions, see [Regular Expressions](#) in the *ICU User Guide*.

### Format file specifications

The format file describes the patterns that the agent looks for to match events in the monitored logs. The format file consists of one or more format specifications.

You can change the format file while an agent instance is running. The file is read by the agent when it starts, and is monitored for changes to its time stamp every 60 seconds thereafter. If the time stamp of the file changes, the agent reinitializes its configuration dynamically, without requiring a restart. For more information, see [“Changing the agent configuration and format files”](#) on page 444.

To create new patterns to match an event, use the new regular expression syntax that consists of the following parts:

- Format header
- Regular expression
- Slot mappings
- End statement

The format header contains the **REGEX** keyword, which informs the agent that you are using a regular expression to match the pattern in the monitored log.

You assign this regular expression to an event class as shown in the following example:

```
REGEX REExample
```

If you use the special predefined event class `*DISCARD*` as your event class, any log records matching the associated pattern are discarded, and no events are generated for them. For example:

```
REGEX *DISCARD*
```

When a pattern is matched, nothing is written to the unmatched log. The log file status records that are matched include these discarded events.

**Note:** You can assign multiple event definitions to either the same event class or to different event classes. The class name is arbitrary and you can use it to indicate the type of event or to group events in various ways.

After the format header, the format content consists of a regular expression on the first line, followed by mappings. Each mapping is shown on a separate line and these mappings are described in the following example.

All lines that match the regular expressions are selected and sent to the monitoring server as events. The regular expression contains subexpressions. You can use the subexpressions to match specific parts of these lines that are the same to a variable called a *slot* in the Event Integration Facility.

The following monitoring log contains three lines that you might want to monitor:

```
Error: disk failure
Error: out of memory
WARNING: incorrect login
```

For example, you generate an event for a specific error, such as the lines that begin with `Error` and ignore the line that begins with `Warning`. The regular expression must match the lines that begin with `Error` and also include a subexpression. The subexpression is denoted by parentheses and it must match only the input text that you want to assign to the *msg* slot. The following format definition is a simple regular expression with only one subexpression:

```
REGEX REExample
Error: (.*)
msg $1
END
```

Based on this format specification, and the preceding set of log data, the agent generates two events. Both events are assigned the `REExample` event class. In the first event, the `disk failure` value is assigned to the *msg* slot. Also, in the second event, the `out of memory` value is assigned to the *msg* slot. Because the `Warning` line did not match the regular expression, it is ignored and no event is generated.

When you assign the value of `$1` to the *msg* slot, you assign it the value of the first subexpression.

If you have log text that contains the following errors, you might want to assign these error messages to their own event class so that you are informed immediately of a disk failure:

```
Error: disk failure on device /dev/sd0: bad sector
Error: disk failure on device /dev/sd1: temperature out of range
```

You can include a description of the disk on which the error occurred, and more specifically the disk error in the event.

The following regular expression contains two subexpressions that identify this information:

```
REGEX DiskFailure
Error: disk failure on device (/dev/sd[0-9]):(.*)
device $1 CustomSlot1
msg $2
END
```

You assign these two subexpressions to event slots. The two events that are generated contain the following values:

```
"device=/dev/sd0" and "msg=bad sector"
"device=/dev/sd1" and "msg=temperature out of range"
```

If you use EIF to generate the first event, it displays as shown in the following example:

```
DiskError;device='/dev/sd0';msg='bad sector';END
```

If the event is sent to the Cloud APM server, the slot that is named *msg* is assigned to the Performance Management agent attribute with the same name. But the *device* slot has no predefined attribute.

If you need to see the value that is assigned to *device* directly on the Cloud APM console, or write thresholds against it, you must assign it to a Performance Management attribute.

The OS agent includes the following 13 predefined attributes:

- Ten string type attributes that range from *CustomSlot1* to *CustomSlot10*
- Three integer type attributes that range from *CustomInteger1* to *CustomInteger3*

Using these attribute names in the format file populates Performance Management attributes with the same name. Using these attributes does not affect the content of the EIF event sent directly to OMNIBus.

**Note:** The `CustomSlot` and `CustomInteger` attribute names are case-sensitive, so you must enter the names exactly as shown.

You assign a slot from the event definition to one of these custom Performance Management attributes in the format file.

You assign the *device* slot to the Performance Management string type attribute called *CustomSlot1* as shown in the following example:

```
REGEX DiskFailure
Error: disk failure on device (/dev/sd[0-9]):(.*)
device $1 CustomSlot1
msg $2
END
```

When the event is displayed in the Application Performance Dashboard, the value that is assigned to the *device* slot is assigned to the Performance Management `CustomSlot1` attribute. You view this value in the Cloud APM console or use it to define thresholds. You can assign any slot in the event definition to any of the 10 custom agent attributes in the same manner, by using "`CustomSlotn`", where *n* is a number from 1 - 10, next to the slot definition.

In this example, the first subexpression is defined specifically as `(/dev/sd[0-9])`, but the second subexpression is defined generally as `(.*)`. In defining the regular expression as specifically as possible, you improve performance. Therefore, if you enter a search for an error on a device that does not match the specific error message that is defined here, the search procedure stops immediately when the error is not found. Time is not wasted looking for a match.

The `END` keyword completes the format specification. The format header, regular expression, and the `END` keyword must each begin on a new line, as shown in the following example:

```
REGEX REExample
Error:
msg $1
END <EOL>
<EOF>
```

**Note:** For the last format in the file, you must insert a new line after the `END` keyword as shown in the example. Otherwise, you get a parsing error.

*CustomInteger1* to *CustomInteger3* are 64-bit custom integer attributes. You can use them in the same manner as the string type `CustomSlot` attributes. You can use these attributes to map individual slots, or subexpressions, from the log file to individual Cloud APM attributes. Because these attributes are numeric, you can use arithmetic comparisons on them, such as `<` and `>`, which is not possible with the string attributes.

**Note:** Although these values are evaluated as integers by the Cloud APM server, for EIF purposes and within the format file, they are still treated as strings. For example, to use an integer slot in a `PRINTF` statement, you still identify it with `"%s"`, not `"%d"`.

The following example illustrates the use of a custom integer attribute. Suppose that a periodic UNIX syslog message is received that reports the percentage of a file system that is free, such as the following hypothetical log record:

```
Oct 24 11:05:10 jimmy fschecker[2165]: Filesystem /usr is 97% full.
```

You can use the following statement in the format file to check for the percentage of the file system that is free:

```
REGEX FileSystemUsage
^([A-Z][a-z]{2}) ([0-9][0-9]) ([0-9]{2}:[0-9]{2}:[0-9]{2}) (.*) (.?):
Filesystem (.*) is ([0-9]+)% full\.$
Month $1 CustomSlot1
Date $2 CustomSlot2
Time $3 CustomSlot3
```

```
Host      $4 CustomSlot4
Service  $5 CustomSlot5
Filesystem $6 CustomSlot6
PctFull  $7 CustomInteger1
msg      PRINTF("%s: %% full", Filesystem, PctFull)
END
```

**Note:** In the preceding statement, everything between the ^ and \$ symbols on the second and third lines must be on a single line.

Because you might have other events that put values in *CustomInteger1*, you can avoid confusing the different event types by using the value of the *Class* attribute to limit its effect to the correct type of events. For example, the following threshold formula causes the threshold to fire only when an event of the *FileSystemUsage* event class has a value greater than or equal to 95 in *CustomInteger1*:

```
( Class == 'FileSystemUsage' AND CustomInteger1 >= 95)
```

A different event can then use *CustomInteger1* for a different purpose and not trigger this threshold accidentally.

In summary, you can now write a threshold in Performance Management that uses arithmetic operators on the *CustomInteger* attributes, which is not possible with the *CustomSlots* attributes.

**Note:** If you map non-integer data to the *CustomInteger* attributes, the resulting value might be zero or some unexpected value.

### Changing the agent configuration and format files

The OS agent reads its configuration (.conf) and format (.fmt) files when it starts, and monitors their time stamp every 60 seconds thereafter.

If the time stamp of the configuration or format file changes, the agent reinitializes its configuration dynamically, without requiring a restart. During reinitialization, monitoring is interrupted momentarily. When monitoring resumes, the agent must determine the position in the monitored logs from which to restart. As a result, the agent behaves in the same way as a full stop and restart.

**Note:** Agent reinitialization after a configuration or format file change resets information in the Log File RegEx Statistics, Log File Status, and Log File Event attribute groups.

By default, the agent starts monitoring from the end of the file, when the reinitialization completes. This starting position can cause events that occurred during the interruption of monitoring to be missed. To ensure that such events are picked up when monitoring resumes, use the `NumEventsToCatchUp=-1` setting.

Setting `NumEventsToCatchUp=-1` causes a position file to be maintained. The position file is updated each time that the agent reads the log file. The update saves the position of the agent in the log file, in case of an agent restart. Maintaining the position file has a small performance impact, so maintain this file only if required. For more information about `NumEventsToCatchUp`, see [“Configuration file”](#) on page 432.

**Note:** Some configuration values are not present in the configuration file and are set during initial configuration. If you change these values, you must restart the agent.

### Inheritance

A format file uses inheritance to derive slot definitions from a previously defined format specification.

Use the `FOLLOWS` relationship to build specific format specifications from generic format specifications by using inheritance.

First, you define a base class and call it `DiskFailure`, for example, as shown here:

```
REGEX DiskFailure
Disk Failure on device (.*
device $1 CustomSlot1
END
```



This regular expression matches the `Disk Failure on device/dev/sd0` errors in the monitoring log so that the `/dev/sd0` value is assigned to the `device` slot.

However, you can also see an extended version of this error message reported in the monitoring log.

For example, you might see a `Disk Failure on device /dev/sd0, error code: 13` error message.

This error message is matched to a slot as shown in the following example:

```
REGEX DiskFailureError FOLLOWS DiskFailure
Disk Failure on device (.*), error code: ([0-9]*)
errcode $2 CustomSlot2
END
```

Now, the event includes the `device` slot and the `errcode` slot. Because the `DiskFailure` event class defined a slot for the device name already, you allow the subclass to inherit that slot, and this inheritance saves you from declaring it a second time. The slot is defined as `$1` so the first subexpression in the regular expression is assigned to that slot.

However, the `DiskFailureError` class also defines a second subexpression. You can assign this subexpression to a new slot called `errcode` and define it as `$2` to refer to the second subexpression in the regular expression. This type of assignment is shown in the previous example that displays the log text.

The event now contains the `device` slot that is assigned the `/dev/sd0` value and the `errcode` slot that is assigned a value of `13`. `CustomSlot1` is assigned the device, and `CustomSlot2` is assigned the error code.

Performance Management custom attribute mappings are also inherited. For more information about Performance Management custom attribute mappings, see [“Format file specifications” on page 441](#).

### Multi-line

Use the multi-line syntax to match records that span more than one line to patterns in the log that you are monitoring.

Specify the `\n` new line character as part of the regular expression to indicate where the line breaks occur in the monitoring log. See this type of syntax in the following example:

```
REGEX REMultiLine
Line1:(.*)\nLine2(.*)
msg $1
second_msg $2
END
```

**Note:** **Windows** Specify a `\r\n` carriage return and new line combination.

If the following error messages are reported in the log text, the `REMultiLine` event is created:

```
Line1: An error occurred
Line2: The error was "disk error"
```

The `msg` slot is assigned the value of `An error occurred` and the `second_msg` slot is assigned the value of `The error was "disk error"`.

### Mappings

The OS agent uses mappings to determine the event class for a system log message. The agent determines the event class by matching the message to a pattern in the format file.

The agent converts log messages to event class instances that contain attribute `name=value` pairs. The event is then sent to the event server.

The agent determines the event class for a system log message at the source. The agent determines the event class by matching a system log message to a pattern in the format file. After you use this matching procedure to determine a class, you must assign values to the attributes.

Attribute values come from various sources, such as:

- Default values that are provided by the agent
- Log text that matches specific subexpressions in regular expressions

A map statement is included in the format file and consists of the following syntax:

```
name    value CustomSlotn
```

Here, you specify any identifier to describe the name of a slot (also known as a variable, attribute, or value identifier). Then, you specify a value to assign to this slot by applying any of the values that are described in [“Value specifiers” on page 446](#).

Use custom slots to view data in the Performance Management console and to define thresholds. When you create thresholds, all custom slot values are strings. Custom slots are also required for duplicate detection to work because you must identify the slots that are used to determine duplicates. For more information about filtering events, see [Chapter 24, “Event filtering and summarization,” on page 1425](#). `msg` is a special slot name, with its own attribute in the event table. You do not need to use a custom slot for the `msg`.

You can limit the scope of a slot so that it exists only within the format definition. When you define the slot, you precede the slot name with a dash, for example:

```
-name    value
```

Any slot that you define in this way is not included in the final event. However, you can reference the slot elsewhere in the format definition, specifically within a PRINTF statement. In the `REGenericSyslog` example that follows, the `service` slot is not included if you generate but you can reference it in the PRINTF statement. It retains the same value that was applied to the original slot when it was defined without the dash. By using this procedure, you can use temporary variables from the format definition that are not included in the final event. For example, you can define an event class, `REGenericSyslog`, to match generic UNIX syslog events in the following way:

```
REGEX REGenericSyslog
^[[A-Z][a-z]{2}} ([ 0-9][0-9]) ([0-9]{2}:[0-9]{2}:[0-9]{2}) (.*) (.*?): (.*)$
month $1
date $2
time $3
host $4
-service $5
msg $6
syslog_msg PRINTF("service %s reports %s", service, msg)
END
```

### *Value specifiers*

The mappings in a format specification assign values to attributes.

The mapping part of a format specification consists of the following types of value specifiers:

- `$i`
- String constant
- PRINTF statement

### **\$i**

The `i` indicates the position of a subexpression in a format string. Each subexpression is numbered from 1 to the maximum number of subexpressions in the format string.

The value of a `$i` value specifier (also known as a variable, slot, or attribute) is the portion of the system log message that is matched by the corresponding subexpression.

In the following example, the log agent translates any log message from the UNIX syslog facility into a syslog event with values assigned to it:

```
REGEX REGenericSyslog
^[[A-Z][a-z]{2}} ([ 0-9][0-9]) ([0-9]{2}:[0-9]{2}:[0-9]{2})
```

```
(.*)? (.*)?: (.*)$
month $1
date $2
time $3
host $4
service $5
msg $6
END
```

Each subexpression numbered from \$1 to \$6 matches an item in parentheses in the regular expression.

Therefore, the following syslog event:

```
Apr 6 10:03:20 jimmy syslogd 1.4.1: restart.
```

is assigned the following values:

```
month=Apr
date=6
time=10:03:20
host=jimmy
service=syslogd 1.4.1
msg=restart.
```

For example, in the syslog event, the 10:03:20 value matches the third item in parentheses in the regular expression, so the value is assigned to the \$3 time value. Similarly, the jimmy value matches the fourth item in parentheses in the regular expression, so the value is assigned to the \$4 host value.

### string constant

The string constant declares that the value of the attribute is the specified string. If the attribute value is a single constant without any spaces, you specify it without surrounding double quotation marks (" ") as shown in the following example:

```
severity WARNING
```

Otherwise, if there are spaces in the attribute value, double quotation marks must be used as shown in the following example:

```
component "Web Server"
```

### PRINTF statement

The PRINTF statement creates more complex attribute values from other attribute values. The PRINTF statement consists of the keyword PRINTF followed by a printf() C-style format string and one or more attribute names.

The format string supports only the %s component specifier. The values of the attributes that are used in the PRINTF statement must be derived from either a \$i value specification or a constant string value specification (you cannot derive them from another PRINTF statement).

Use the value of the argument attributes to compose a new constant string according to the format string. This new constant string becomes the value of the attribute.

Based on the previous example where you defined the REGenericSyslog base class, and the service and msg slots, you can define an attribute called syslog\_msg by using the PRINTF keyword.

```
syslog_msg PRINTF("service %s reports %s", service, msg)
```

If the following log message is reported:

```
Apr 6 10:03:20 jimmy syslogd 1.4.1: restart.
```

a new constant string is composed that contains the attribute values from the format string:

```
syslog_msg="service syslogd 1.4.1 reports restart."
```

## Keywords

In the format file, use keywords to assign values that expand at run time.

The following keywords expand at run time:

- DEFAULT
- FILENAME
- LABEL
- REGEX

## DEFAULT

Use the DEFAULT keyword to assign a DEFAULT value to a specific slot or attribute. The OS agent assigns an internal default value to slots that are described in the following table:

| Slots             | Description                                                                                                                                                                                                                                                                                                              |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>hostname</i>   | <i>hostname</i> is the short host name of the system where the agent is running. It does not include the domain name of the system.                                                                                                                                                                                      |
| <i>origin</i>     | <i>origin</i> is the IP address of the system where the agent is running.                                                                                                                                                                                                                                                |
| <i>fqhostname</i> | <i>fqhostname</i> is the fully qualified host name of the system where the agent is running. It includes the domain name of the system.                                                                                                                                                                                  |
| <i>RemoteHost</i> | When an event originates on the local system, this attribute is empty. If an event originates on a remote system, <i>RemoteHost</i> contains a string of the form <i>user@host:port</i> , which indicates the remote host name on which the event occurred, and the user and port on that host that are used to connect. |

The value that is assigned to *fqhostname* is influenced by the following FQDomain (optional) settings in the `.conf` file:

- If you set FQDomain to yes, the agent determines the system domain name itself.
- If you do not set a value for FQDomain or if you set the value to no, the agent does not set a domain name, and the *fqhostname* attribute is assigned a blank string.
- If you set FQDomain so that it does not contain a yes or no value, the domain name is accepted as the value and it is appended to the host name.

In the following example, the format definition contains three attributes or slots:

- *hostname* DEFAULT
- *origin* DEFAULT
- *fqhostname* DEFAULT

If you set the FQDomain to yes in the `.conf` file and you run it on a computer with the following properties:

- *hostname*: myhost
- *IP address*: 192.168.1.100
- *domainname*: mycompany.com

an event is created and the three slots are assigned the following values:

```
"hostname=myhost", "origin=192.168.1.100", "fqhostname=myhost.mycompany.com"
```

## FILENAME

The FILENAME keyword indicates the fully qualified file name (including the path) of the log file that contains the message. If you use a single agent to monitor multiple log files and you need to identify the source of the event, use this keyword to populate an event attribute with the file name. If the message comes from the system log, mapping is set to EventLog for Windows OS agents and SysLogD for UNIX OS agents.

**Note:** The path includes an attribute for this keyword.

## LABEL

The LABEL keyword specifies the host name of the system where the agent is running.

## REGEX

The REGEX keyword expands to the regular expression that matched the message and caused the event.

### Maximum message length

This value is the maximum message length that the OS agent can receive without truncating the message.

The maximum message length is different for Performance Management and Tivoli Netcool/OMNIBus.

### Performance Management

For events sent to Performance Management, the msg attribute is limited to 2048 bytes. Messages that are greater in length are truncated.

### Tivoli Netcool/OMNIBus

For events sent through the Probe for Tivoli EIF to Netcool/OMNIBus, the total size of the event, including the class name and all slots and their values cannot exceed 4096 bytes. For example, in the following sample EIF event, ;END does not count against the 4096-byte limit. However, everything else does count against the limit, including the syntactic elements such as the semicolons, quotation marks, and equal signs.

```
Class;attr1=value1;attr2=value2;msg='Hello, world';END
```

## Remote log file monitoring: Encrypting a password or pass phrase

For increased security, you can encrypt passwords and pass phrases that are transmitted to remote systems when you use Remote log file monitoring.

### About this task

The encrypted password and pass phrases are stored in the configuration (.conf) file. For more information about the configuration file, see [“Configuration file” on page 432](#).

### Procedure

- Run the **itmpwdsnmp** command and supply the password or pass phrase that is to be encrypted:
  - **Linux** | **UNIX** The command is run from the Cloud APM installation directory. The default installation path is `opt/ibm/apm/agent` and `install_dir` is where you installed the agent.
  - **Windows** The default installation path is `C:\IBM\APM`.

**Linux** Example of the command when it is run on a Linux system:

```
$ export install_dir=/opt/ibm/apm/agent/bin
$ /opt/ibm/apm/agent/bin
```

**Enter string to be encrypted:**

mypassword

**Confirm string:**

```
mypassword
```

```
{AES256:keyfile:a}Z7BS23aupYqw1Xb1Gh+weg==  
$
```

In the example, the entire output from the `{AES256:keyfile:a}Z7BS23aupYqw1Xb1Gh+weg==` command is used to set **SshPassword** in the agent configuration file. The `{AES256:keyfile:a}` prefix tells the agent that the password is encrypted.

To encrypt a pass phrase for a private key file, follow the same procedure.

## Configuring Linux OS Agent file system data collection

The Monitoring Agent for Linux OS is configured automatically. However, you can configure the behavior for file system data collection.

The Monitoring Agent for Linux OS has default behavior for file system data collection.

The default behavior is to monitor file systems from the `/etc/fstab` only. An environment variable `KBB_SHOW_MTAB_FS` is defined in the `lz.environment` file to control the file system data collection behavior. If you want to monitor all file systems (listed in `/etc/fstab` and `/etc/mtab`), you can set `KBB_SHOW_MTAB_FS=true`.

### KBB\_SHOW\_MTAB\_FS

This variable is available in the `install_dir/config/lz.environment` file. The default value is `false` and defines the agent to monitor file systems from the `/etc/fstab` only. If you want to monitor all file systems (listed in `/etc/fstab` and `/etc/mtab`), change the value to `true`. For example, `KBB_SHOW_MTAB_FS=true`.

## Configuring PostgreSQL monitoring

Configure the Monitoring Agent for PostgreSQL so that the agent can collect data from the PostgreSQL database that is being monitored.

### Before you begin

You must install the PostgreSQL JDBC driver before you install the PostgreSQL agent. The path to PostgreSQL JDBC driver is required at the time of agent configuration.

JDBC type 4 driver is the most recent version and hence must be preferred. User can install the subtype of JDBC 4 version according to the JDK version the agent uses. For more information about mapping JDBC version to JDK version, see [PostgreSQL JDBC Driver](#).

A few of the attributes that are collected by the agent rely on the `pg_stat_statements` extension. To add `pg_stat_statements`, first install the `postgresql-contrib` package. You must modify the `postgresql.conf` configuration file for the PostgreSQL server to load the `pg_stat_statements` extension.

1. Open the `postgresql.conf` file in a text editor and update the `shared_preload_libraries` line:

```
shared_preload_libraries = 'pg_stat_statements'  
pg_stat_statements.track_utility = false  
listen_addresses='<host_ip_address>'
```

Where the `<host_ip_address>` is the IP address of the Virtual machine where PostgreSQL agent is installed. You can modify the value of `<host_ip_address>` parameter as `*`, which means that it can accept IP addresses of all hosts.

These changes are required to monitor SQL statements, except utility commands.

**Note:** The status of `pg_stat_statements.track_utility` is set or modified by a superuser only.

2. Restart the PostgreSQL server after you update and save the `postgresql.conf`.
3. Run the following SQL command by using `psql` that must be connected to the same database that would be provided later in the agent configuration for JDBC connectivity:

```
create extension pg_stat_statements;  
select pg_stat_statements_reset();
```

**Note:** The command `create extension` and function `pg_stat_statements_reset()` are run by a superuser only.

The view `pg_stat_statements` needs to be enabled for specific database, for more details refer [pg\\_stat\\_statements](#).

The `pg_hba.conf` file contains authentication settings of PostgreSQL database. When the `auth-method` parameter value is set to `ident` in the `pg_hba.conf` file, the PostgreSQL agent cannot connect to the PostgreSQL database. Ensure that the authentication settings for the `auth-method` parameter are correct. For example, you can set these values for `auth-method` parameter: `md5`, `trust`, or `password`.

### About this task

The PostgreSQL agent is a multiple instance agent. You must create the first instance and start the agent manually. The managed system name includes the instance name that you specify, for example, `instance_name: host_name`. The managed system name is limited to 32 characters. The instance name that you specify is limited to 28 characters, minus the length of your host name. For example, if you specify `PostgreSQL2` as your instance name, your managed system name is `PostgreSQL2: hostname`.

**Important:** If you specify a long instance name, the managed system name is truncated and the host name is not displayed completely.

The product version and the agent version often differ. The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For information about the agent version list and what's new for each version, see the ["Change history"](#) on page 52.

## Configuring the agent on Windows systems

You can use the Application Performance Management window to configure the agent on Windows systems.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Monitoring Agent for PostgreSQL**, and then click **Configure agent**.
3. In the **Enter a unique instance name** field, type the agent instance name and click **OK**.
4. In the **Monitoring Agent for PostgreSQL** window, complete the following steps:
  - a. In the **IP Address** field, enter the IP address for PostgreSQL server that you want to monitor remotely. If the agent is installed on the server to be monitored, retain the default value.

#### Note:

For remote monitoring, the data for **Current CPU used(%)** and **Physical memory used(MB)** is displayed as **N/A** on the dashboard.

- b. In the **JDBC database name** field, enter a database name to change the default database name of `postgres`.
- c. In the **JDBC user name** field, enter a user name to change the default name of `postgres`.
- d. In the **JDBC password** field, enter the JDBC user password.
- e. In the **Confirm JDBC password** field, re-enter the password.
- f. In the **JDBC port number** field, enter a port number to change the default port number of 5432.
- g. In the **JDBC JAR file** field, enter the path for the PostgreSQL connector for the Java JAR file and click **Next**.

- h. In the **Java trace level** field, enter the trace level according to the IBM support instructions. The default trace level is `ERROR`.
  - i. Click **OK**. The agent instance is displayed in the IBM Performance Management window.
5. Right-click the **Monitoring Agent for PostgreSQL** instance, and click **Start**.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent on Linux systems

To configure the agent on Linux operating systems, you must run the script and respond to prompts.

### Procedure

1. On the command line, enter the following command:  
`install_dir/bin/postgresql-agent.sh config instance_name`
2. When you are prompted to edit the agent for PostgreSQL settings, enter `1` to continue.
3. When you are prompted to enter a value for the following parameters, press Enter to accept the default value or specify a different value and press Enter:

- PostgreSQL server IP address

#### Note:

Enter the IP address of a PostgreSQL server that you want to monitor remotely. If the agent is installed on the server to be monitored, retain the default value.

For remote monitoring, the data for **Current CPU used(%)** and **Physical memory used(MB)** is displayed as **N/A** on the dashboard.

- JDBC database name
- JDBC user name
- JDBC password
- JDBC port number
- JDBC JAR file

**Important:** The version of the JDBC JAR file must be same as the version of the PostgreSQL server that is monitored.

4. When you are prompted to enter a value for the `Java trace level` parameter, press enter to accept the default value or specify the trace level according to the IBM support instructions.
5. Run the following command to start the agent:

```
install_dir/bin/postgresql-agent.sh start instance_name
```

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [“Starting the Cloud App Management UI” on page 176](#).



## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

You can use the silent response file to configure the PostgreSQL agent on Linux and Windows systems. After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

- To configure the agent by editing the silent response file and running the script without responding to prompts, complete the following steps:
  1. In a text editor, open the silent response file that is available in this path: `install_dir/samples/postgresql_silent_config.txt`  
Where `install_dir` is the installation directory of PostgreSQL agent. The default installation directory is `/opt/ibm/apm/agent`.
  2. To edit the silent configuration file, complete the following steps:
    - a. For the **IP Address** parameter, specify the IP address of a PostgreSQL server that you want to monitor remotely. If the agent is installed on the server to be monitored, retain the default value.  
**Note:** For remote monitoring, the data for **Current CPU used(%)** and **Physical memory used(MB)** is displayed as **N/A** on the dashboard.
    - b. For the **JDBC database name** parameter, specify a database name to change the default database name of postgres.
    - c. For the **JDBC user name** parameter, specify a user name to change the default name of postgres.
    - d. For the **JDBC password** parameter, enter the JDBC user password.
    - e. For the **JDBC port number** parameter, specify a port number to change the default port number of 5432.
    - f. For the **JDBC JAR file** parameter, specify the path for the PostgreSQL connector for the Java JAR file if the default path is incorrect. The default path of the Java JAR file is:  
`/opt/PostgreSQL/lib/postgresql-9.3-1100.jdbc4.jar`  
**Important:** The version of the JDBC JAR file must be compatible with the version of the PostgreSQL database that is being monitored.
    - g. For the **Java trace level** parameter, specify the trace level according to the IBM support instructions. The default trace level is `Error`.
  3. Save and close the silent response file, and run the following command:

```
install_dir/bin/postgresql-agent.sh config instance_name \  
install_dir/samples/postgresql_silent_config.txt
```

Where `instance_name` is the name that you want to give to the instance.

4. To start the agent, enter the following command:

```
install_dir/bin/postgresql-agent.sh start instance_name
```

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring RabbitMQ monitoring

The Monitoring Agent for RabbitMQ monitors the health and performance of the RabbitMQ cluster resources, such as the nodes, queues, and channels of the cluster. You must configure the RabbitMQ agent so that the agent can collect the RabbitMQ data.

### Before you begin

- Review the hardware and software prerequisites.
- Ensure that the RabbitMQ user, who connects to the node, has read permission and either the monitoring, administrator, or management tag is enabled for this user.
- Ensure that the RabbitMQ management plugin is enabled on all nodes of the cluster, because if one node of the cluster fails, the RabbitMQ agent connects to a peer node that is available in the cluster.

### About this task

The product version and the agent version often differ. The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [“Using agent commands” on page 226](#). For detailed information about the agent version list and what's new for each version, see the [“Change history” on page 52](#).

The RabbitMQ agent is a multiple instance agent. You must create the first instance, and start the agent manually.

## Configuring the agent on Windows systems

You can use the **IBM Performance Management** window to configure the agent on Windows systems.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Monitoring Agent for RabbitMQ**, and then click **Configure agent**.

**Remember:** After you configure the agent for the first time, the **Configure agent** option is disabled. To configure the agent again, click **Reconfigure**.

3. In the **Enter a unique instance name** field, type the agent instance name and click **OK**.
4. In the **Monitoring Agent for RabbitMQ** window, specify values for the configuration parameters, and then click **Next**.

For information about the configuration parameters, see [“Configuration parameters for the agent” on page 456](#).

5. Right-click **Monitoring Agent for RabbitMQ** instance, and click **Start**.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent on Linux systems

To configure the agent on Linux operating systems, you must run the script and respond to prompts.

### Procedure

1. On the command line, enter the following command:

```
install_dir/bin/rabbitmq.sh config instance_name
```

Where *instance\_name* is the name that you want to give to the instance.

2. When you are prompted to provide a value for the following parameters, press Enter to accept the default value, or specify a value and then press Enter:

- IP Address
- User Name
- Password
- Port Number
- Java home
- Java trace level

For information about the configuration parameters, see [“Configuration parameters for the agent” on page 456](#).

3. Run the following command to start the agent:

```
install_dir/bin/rabbitmq.sh start instance_name
```

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### About this task

You can use the silent response file to configure the RabbitMQ agent on Linux and Windows systems. After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

1. Open the silent response file from the following location:

```
install_dir\samples
```

2. In the `rabbitmq_silent_config.txt` file, specify values for all mandatory parameters. You can also modify the default values of other parameters.

For information about the configuration parameters, see [“Configuration parameters for the agent” on page 456](#).

3. Save the response file, and run the following command:

```
Linux | UNIX install_dir/bin/rabbitmq-agent.sh config install_dir/samples/rabbitmq_silent_config.txt
```

```
Windows install_dir/bin/rabbitmq-agent.bat config install_dir/samples/rabbitmq_silent_config.txt
```

4. Start the agent using the following command:

```
Linux | UNIX Run the following command: install_dir\bin\rabbitmq-agent.sh start
```

```
Windows Right-click Monitoring Agent for RabbitMQ and then click Start.
```

## What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuration parameters for the agent

When you configure the RabbitMQ agent, you can change the default values of the parameters, such as the instance name and the SSL validation certificates.

The following table contains detailed descriptions of the configuration parameters for the RabbitMQ agent.

| Parameter name   | Description                                                                                                                                                                                                              | Mandatory field |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| IP Address       | The IP address of the node where the RabbitMQ application is installed.                                                                                                                                                  | Yes             |
| Username         | The user name of the RabbitMQ user.                                                                                                                                                                                      | Yes             |
| Password         | The password to connect to the RabbitMQ management user interface.                                                                                                                                                       | Yes             |
| Confirm Password | The same password that you entered in the <b>Password</b> field.                                                                                                                                                         | Yes             |
| Port Number      | The port number where the RabbitMQ management plugin is enabled. Use the default port number 15672, or specify another port number.                                                                                      | No              |
| Java home        | The path where the java plugin is installed. Use the default path C:\Program Files\IBM\Java50, or the directory path where java plugin is installed.                                                                     | No              |
| Java trace level | The trace level of the Java provider. The valid trace level values are as follows: <ul style="list-style-type: none"><li>• OFF</li><li>• ERROR</li><li>• WARN</li><li>• INFO</li><li>• DEBUG_MAX</li><li>• ALL</li></ul> | No              |

## Configuring SAP monitoring

To monitor a SAP system, the Monitoring Agent for SAP Applications must connect to an application server in the system to be monitored so that the agent can access the Advanced Business Application Programming (ABAP) code that is provided with the product.

### Before you begin

- The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history” on page 52](#).
- Review the hardware and software prerequisites, see [Software Product Compatibility Reports for SAP agent](#)
- The SAP agent does not support non-Unicode SAP systems.

## About this task

The SAP agent is a multiple instance agent. You must create the first instance and start the agent manually.

- To configure the agent on Windows systems, you can use the **IBM Performance Management** window or the silent response file.
  - [“Configuring the agent on Windows systems” on page 457](#)
  - [“Configuring the agent by using the silent response file” on page 459](#)
- To configure the agent on Linux or AIX systems, you can run the script and respond to prompts, or use the silent response file.
  - [“Configuring the agent on Linux or AIX systems” on page 458](#)
  - [“Configuring the agent by using the silent response file” on page 459](#)

After you install the SAP agent, you can import the Advanced Business Application Programming (ABAP) transport on the SAP system to support data collection in the SAP system. For more information, see [“Importing the ABAP transport on the SAP system” on page 462](#).

After you configure the SAP agent, you must verify the agent configuration. For more information, see [“Verifying agent configuration” on page 469](#).

To delete the ABAP transport from the SAP system, you must import delete transport to the SAP system. For more information, see [“Deleting the ABAP transport from the SAP system” on page 468](#).

The new CCMS design is enabled by default. Entry is present in the database table /IBMMON/ITM\_CNGF for **isnewccmsdesign** parameter whose value is set to **YES**.

## Configuring the agent on Windows systems

You can configure the SAP agent on Windows systems by using the **IBM Performance Management** window so that the agent can collect data of the SAP Applications Server that is being monitored.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Template** under the **Task/SubSystem** column, and click **Configure Using Defaults**.

The **Monitoring Agent for SAP Applications** window opens.

3. In the **Enter a unique instance name** field, type an agent instance name and click **OK**.

**Important:** The agent instance name must match the 3-digit system identifier (SID) of the managed SAP Applications Server. For example, if the SID of the managed SAP Applications Server is PS1, enter PS1 as the instance name.

4. Configure the SAP agent in the Application Server mode or the Logon Group mode.

- To configure the SAP agent in the Application Server mode complete the following steps:
  - a. In the **Connection Mode** field, select **Application Server Mode** and click **Next**.
  - b. In the **Specify Application Server Information** area, specify values for the configuration parameters and click **Next**.
  - c. In the **Specify Logon Information to the SAP System** area, specify values for the configuration parameters and click **OK**.

For more information, see [“Configuration parameters of the agent” on page 459](#)

- To configure the SAP agent in the Logon Group mode complete the following steps:
  - a. In the **Connection Mode** field, select **Logon Group Mode** and click **Next**.
  - b. In the **Specify Logon Group Information** area, specify values for the configuration parameters and click **Next**.

- c. In the **Specify Logon Information to the SAP System** area, specify values for the configuration parameters and click **OK**.

For more information, see [“Configuration parameters of the agent” on page 459](#)

**Important:** For the Application Server mode, it is mandatory to configure the Dialog Instance having dispatcher on the SAP system where the message server or ASCS is configured. For the Logon Group mode, it is not mandatory to configure the Dialog Instance having dispatcher on the SAP system where the message server or ASCS is configured.

5. In the **IBM Performance Management** window, right-click the agent instance that you created and click **Start**.

**Important:** If you want to create another instance of the SAP agent, repeat Steps 1 - 6. Use a unique system identifier for each SAP agent instance that you want to create.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [Starting the Cloud App Management UI](#).

## Configuring the agent on Linux or AIX systems

You can configure the SAP agent on Linux or AIX systems so that the agent can collect data of the SAP Applications Server that is being monitored.

### Procedure

1. On the command line, change the path to the agent installation directory.  
For example, `/opt/ibm/apm/agent/bin`
2. Run the following command:

```
./sap-agent.sh config instance_name
```

where, *instance\_name* is the name that you want to give to the instance.

**Important:** The agent instance name must match the 3-digit system identifier (SID) of the managed SAP Applications Server. For example, if the SID of the managed SAP Applications Server is PS1, enter PS1 as the instance name.

3. When the command line displays the following message, type 1 and press Enter. Edit 'Monitoring Agent for SAP Applications' setting? [1=Yes, 2=No]
4. Configure the SAP agent by using the Application Server mode or the Logon Group mode.
  - To configure the SAP agent in the Application Server mode complete the following steps:
    - a. When the command line displays the following message, type 1 and press Enter: Connection Mode [ 1=Application Server Mode, 2=Logon Group Mode ]
    - b. Specify values for the configuration parameters. For information, see [“Configuration parameters of the agent” on page 459](#).
  - To configure the SAP agent in the Logon Group mode complete the following steps:
    - a. When the command line displays the following message, type 2 and press Enter: Connection Mode [ 1=Application Server Mode, 2=Logon Group Mode ]
    - b. Specify values for the configuration parameters. For more information about the configuration parameters, see [“Configuration parameters of the agent” on page 459](#).

**Important:** For the Application Server mode, it is mandatory to configure the Dialog Instance having dispatcher on the SAP system where the message server or ASCS is configured. For the Logon Group mode, it is not mandatory to configure the Dialog Instance having dispatcher on the SAP system where the message server or ASCS is configured.

5. Run the following command to start the SAP agent:

```
./sap-agent.sh start instance_name
```

**Important:** If you want to create another instance of the SAP agent, repeat Steps 1 - 5. Use a unique System Identifier for each SAP agent instance that you create.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [Starting the Cloud App Management UI](#).

## Configuring the agent by using the silent response file

You can configure the SAP agent on Windows, Linux, or AIX systems by using the silent response file.

### Procedure

1. In a text editor, open the `sap_silent_config.txt` file that is available at the `install_dir\samples` path, and specify values for all the configuration parameters.

- For windows systems, the default file path is `C:\IBM\APM\samples`
  - For Linux and AIX systems, the default file path is `/opt/ibm/apm/agent/samples`
- For more information, see [“Configuration parameters of the agent” on page 459](#).

2. Change the file path as follows:

- For Windows systems, the file path is `install_dir\BIN`
- For Linux and AIX systems, the file path is `install_dir\bin`

3. Run the following command.

- On Windows systems:

```
sap-agent.bat config instance_name install_dir\samples\sap_silent_config.txt
```

- On Linux and AIX systems:

```
sap-agent.sh config instance_name install_dir\samples\sap_silent_config.txt
```

**Important:** The agent instance name must match the 3-digit system identifier (SID) of the managed SAP Applications Server. For example, if the SID of the managed SAP Applications Server is PS1, enter PS1 as the instance name.

4. Start the agent.

- On Windows systems, in the **IBM Performance Management** window, right-click the agent instance that you created, and click **Start**.
- On the Linux and AIX systems run the following command:

```
./sap-agent.sh start instance_name
```

**Important:** If you want to create another instance of the SAP agent, repeat Steps 1 - 4. Use a unique System Identifier for each SAP agent instance that you create.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [Starting the Cloud App Management UI](#).

## Configuration parameters of the agent

When you configure the SAP agent, you can change the default value of the parameters, such as the SAP hostname and the SAP system number.

The following table contains detailed descriptions of configuration parameters of the SAP agent.

Table 57. Names and descriptions of configuration parameters of the SAP agent

| Parameter name                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                             | Mandatory field | Examples           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|--------------------|
| SAP Hostname (Primary)          | The host name of the SAP application server to which the agent connects. If your SAP servers communicate over a private LAN, the computers that host the servers have two or more network cards. For the host name, enter a name by which the application server can be reached from external systems, such as the SAPGUI logon. Do not use the private LAN host name. The default value is the host name where the agent is installed. | Yes             | saphost.domain.com |
| SAP System Number (Primary)     | The two-digit SAP system or instance number that is used for connecting to a SAP host server. The default value is 00.                                                                                                                                                                                                                                                                                                                  | Yes             |                    |
| SAP Hostname (Alternate 1)      | The second choice for the host name if the primary host is unavailable.                                                                                                                                                                                                                                                                                                                                                                 | No              |                    |
| SAP System Number (Alternate 1) | The system number for the host name of the first alternate.                                                                                                                                                                                                                                                                                                                                                                             | No              |                    |
| SAP Hostname (Alternate 2)      | The third choice for the host name if both the SAP Hostname (Primary) and SAP Hostname (Alternate 1) hosts are unavailable.                                                                                                                                                                                                                                                                                                             | No              |                    |
| SAP System Number (Alternate 2) | The system number for the host name of the second alternate.                                                                                                                                                                                                                                                                                                                                                                            | No              |                    |
| SAP Client Number               | The SAP client number for the RFC logon to SAP. The default value is 000. If the IBMMON_AGENT user that is generated by ABAP is used, enter the client number that was specified in the transport import. This number is the same as the nnn client number under the profile.                                                                                                                                                           | Yes             |                    |
| SAP User Id                     | The SAP user ID for the RFC logon to SAP. The default value is IBMMON_AGENT, which is the predefined user ID that is created during the import.                                                                                                                                                                                                                                                                                         | Yes             |                    |
| SAP User Password               | Use the default password or specify a different password.                                                                                                                                                                                                                                                                                                                                                                               | Yes             |                    |
| Confirm SAP User Password       | The password that is specified in the <b>SAP User Password</b> field.                                                                                                                                                                                                                                                                                                                                                                   | Yes             |                    |



Table 57. Names and descriptions of configuration parameters of the SAP agent (continued)

| Parameter name          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Mandatory field | Examples |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|----------|
| SAP Language Code       | <p>The language code that indicates the language that the agent uses when it connects to the SAP system. The specified language determines the language in which you see SAP information, such as alert messages, syslog messages, and job log messages.</p> <p>All SAP systems are delivered in English and German. If you require a different language, confirm with your SAP administrator that the language is installed on the SAP system. If you specify a language that is not supported, the agent cannot connect to the SAP system.</p> <p>The following languages and codes are supported:</p> <ul style="list-style-type: none"> <li>• CS - Czech</li> <li>• EN - English</li> <li>• FR - French</li> <li>• DE - German</li> <li>• HU - Hungarian</li> <li>• IT - Italian</li> <li>• ES - Spanish</li> <li>• JA - Japanese</li> <li>• KO - Korean</li> <li>• PL - Polish</li> <li>• PT - Portuguese</li> <li>• RU - Russian</li> <li>• ZH - Chinese</li> <li>• ZF - Traditional Chinese</li> </ul> | Yes             |          |
| RFC Trace               | <p>The Remote Function Call (RFC) trace setting for the <code>SAPTRACE</code> variable. When you select this check box, you activate the RFC tracing and the default value is no RFC tracing. For the command line, 2 = No trace and 1 = Do trace. Because the RFC tracing generates extensive diagnostic information, use it carefully. For more information about the RFC tracing, contact IBM support.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | No              |          |
| SAP Logon Group         | The name of the SAP Server Logon group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Yes             |          |
| SAP Message Server Name | The host name of the SAP message server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Yes             |          |

Table 57. Names and descriptions of configuration parameters of the SAP agent (continued)

| Parameter name                 | Description                                                                                                                                                                                                                                                              | Mandatory field | Examples                                                                                                                                        |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| SAP Message Service            | The name of the service where the SAP message server is located.<br>You must include service names in the following operating system services files: <ul style="list-style-type: none"> <li>• /etc/services</li> <li>• \windows\system32\drivers\etc\services</li> </ul> | Yes             | You might use the message service name sapmsTV1, or the full message service port number 3601.                                                  |
| SAP Route String               | Specify the SAP router string if you want access to the SAP server with a SAP router.                                                                                                                                                                                    | No              | The router string /H/host/H/ must be in the following format: /H/beagle/H/brittany/H/ or /H/amsaix11.tivlab.raleigh.ibm.com/W/tivoli/H/amsaix25 |
| SNC                            | Specify whether you want to enable or disable Secure Network Communications (SNC). Default value is disabled.                                                                                                                                                            | Yes             | <b>sap_conn.sap_snc_mode</b><br>=true or false                                                                                                  |
| SNC Security Level             | The security level of SNC.                                                                                                                                                                                                                                               | Yes             | <b>sap_snc_mode1.sap_snc_qop</b> =QOP value.<br>Default value is 8.                                                                             |
| Client or Agent SNC Name       | The SNC name of the client or agent.                                                                                                                                                                                                                                     | Yes             | <b>sap_snc_mode1.sap_snc_client</b> =<br>Client SNC Name                                                                                        |
| Partner or SAP Server SNC Name | The SNC name of the partner or SAP Server.                                                                                                                                                                                                                               | Yes             | <b>sap_snc_mode1.sap_snc_server</b> =<br>Server SNC Name                                                                                        |
| SAP Cryptolibrary Path         | The path of SAP Cryptolibrary.                                                                                                                                                                                                                                           | Yes             | <b>sap_snc_mode1.sap_snc_library</b> =<br>Crypto library path                                                                                   |

## Importing the ABAP transport on the SAP system

You can install one SAP agent for each SAP system where you import the Advanced Business Application Programming (ABAP) transport request to support data collection in the SAP system.

### Before you begin

Before you import ABAP transport on the SAP system, ensure that the following prerequisites are met:

- To import the product transport request, R3trans Version 01.07.04, or later is required because the Dynpro and export and import tables are incompatible. The basic operation of the agent is not affected

by the Dynpro or export and import incompatibility issues, only the SAP configuration windows are affected.

- You must ensure that you import the SAP agent transport on the client where the MAI configuration is available to monitor the Solution Manager System. To view features of the PI system, import the SAP agent transport on the PI system on a client where PI configuration is available.
- To view data in the group widgets that are under SLM subnode, you must complete the MAI configurations for PI and Solution Manager. You must also configure business process monitoring so that you can view data in the BPM Alerts group widget. To view data for the Latest Critical and High Priority Alerts group widget, make the following configurations:
  - In Solution Manager 7.1, run SOLMAN\_SETUP transaction and select **System Monitoring**, activate or enable the third-party component, and add **Implementation: BADI Definition for Alert Reactions** and third-party connector.
  - Set the scope filter to **All Alerts and Metrics**.
  - Ensure that the implementation state is **Active**.

For more information, see the following Online Service System (OSS) Notes, which include a list of required SAP service pack levels:

- OSS Note 454321
  - OSS Note 330267
  - OSS Note 743155
- To monitor the SAP systems, the SAP agent needs the SAP statistics data. On SAP 7.0 systems, you must set the SAP system time to match the time for the operating system so that SAP statistics are collected with the correct time stamps. Similarly, update the SAP system time for the SAP agent so that the agent can collect data. For more information about this issue, see SAP Note 926290.

### About this task

For more information, see [“Importing the SAP transport”](#) on page 464.

### MAI Alert related prerequisites for importing the ABAP transport

You must verify the Monitoring and Alerting Infrastructure (MAI) Alert related prerequisites before you import the ABAP transport.

### Configuration settings in the transport .prop file

When you use the new MAI Alert fetching mechanism that includes fetching MAI Alerts without configuring email notification settings and without BAdi implementation, then you must modify the following configuration setting in the `transport .prop` file.

Add the `SPLEVEL=X` line, where X is the support pack (SP) level of the Solution Manager system. For example, if the System ID is S10 and the support pack level is 13, then add `SPLEVEL=13`.

**Important:** For the SAP system with SP level 10, or later, the value of the Technical Name (MEA) attribute is not populated on the Latest MAI Alerts with Rating 'Red' group widget in the SAP Solution Manager Dashboard when the MAI Alerts are fetched without configuring email notification in the SAP Solution Manager and without BAdi implementation. It gets populated when the MAI Alerts are fetched by configuring email notification in the SAP Solution Manager and BAdi implementation.

### Determination of old and new mechanism for fetching MAI Alerts based on the Solution Manager Support Pack (SP) Level

#### Old MAI Alert fetching mechanism

This mechanism is based on configuring email notification settings and the `/IBMMON/ITM_IMPL_ALRTINBX` BAdi implementation with the `IF_ALERT_DYN_COFIGURATION` interface to collect MAI Alerts and send them to the SAP agent.

### New MAI Alert fetching mechanism

This mechanism is based on fetching MAI Alerts without configuring email notification settings and without the /IBMMON/ITM\_IMPL\_ALRTINBX BAdi implementation with the IF\_ALERT\_DYN\_COFIGURATION interface.

You can use the following table to understand the usage of the transport . prop file and its dependency on the configuration of email notification settings.

*Table 58. Usage of transport.prop file and its dependencies*

| SAP system SP Level | transport.prop settings   |                           | Configuration of email notification settings | MAI Alert mechanism to be used                                                                     |
|---------------------|---------------------------|---------------------------|----------------------------------------------|----------------------------------------------------------------------------------------------------|
|                     | MAI_CONFIGURED            | Solution Manager SP level |                                              |                                                                                                    |
| Any                 | No or file does not exist | Not Applicable            | Configured or not configured                 | The SLM subnode does not appear instead the SOL subnode appears.                                   |
| SP 6 - 9            | Yes                       | Mentioned                 | Configured                                   | Old mechanism                                                                                      |
| SP 6 - 9            | Yes                       | Not mentioned             | Configured                                   | Old mechanism                                                                                      |
| SP 6 - 9            | Yes                       | Not mentioned             | Not configured                               | Old mechanism does not work because the configuration of email notification settings is mandatory. |
| SP 6 - 9            | Yes                       | Mentioned                 | Not configured                               | Old mechanism does not work because the configuration of email notification settings is mandatory. |
| SP 10, or later     | Yes                       | Mentioned                 | Configured                                   | New mechanism                                                                                      |
| SP 10, or later     | Yes                       | Mentioned                 | Not configured                               | New mechanism                                                                                      |
| SP 10, or later     | Yes                       | Not mentioned             | Configured                                   | Old mechanism                                                                                      |
| SP 10, or later     | Yes                       | Not mentioned             | Not configured                               | Old mechanism does not work because the configuration of email notification settings is mandatory. |

### Importing the SAP transport

The SAP agent provides a set of Advanced Business Application Programming (ABAP) routines to support data collection in the SAP system. This ABAP code is delivered as an SAP transport that must be installed on each SAP system that is to be monitored. Your SAP administrator installs the transport.

### About this task

The **ZITM\_610AUTH** authorization profile and **ZITM\_610AUT** authorization role are valid until the 6.1 release only. From release 6.2 or later, the **/IBMMON/AUTH** authorization profile is used. To protect

against unauthorized use, the ABAP code that is installed in the SAP system is not visible from within the SAP system. In addition, this code cannot be modified or generated. You must obtain the support for this code from the IBM software support website.

In addition to installing ABAP code, the transport also installs translated language text elements to provide multicultural support for SAP transport text elements.

**Important:** Before you import the transport into the SAP system, you must not start the SAP agent instance that is configured to monitor the SAP system.

When you import the SAP transport, users get implicitly defined in the SAP system. You can import the SAP transport into the SAP system as follows.

## Procedure

To import the SAP transport into the SAP system follow these steps.

1. Copy the IBM Tivoli Monitoring transport file from the following paths on the computer where the agent is installed.
  - For Windows systems, the file path is *install\_dir\TMAITM6\_x64\ABAP*
  - For Linux and AIX systems, the file path is *install\_dir/intrp/sa/ABAP*, here *intrp* must be **1x8266** or **aix526**.
2. Copy the following transport files from the paths that are mentioned in step 1 into the SAP environment:
  - K711\_00xxxU.ITM and R711\_00xxxU.ITM files are Unicode versions of the transport. They contain the SAP agent ABAP code and Unicode support for text strings for Latin code pages and double-byte code pages.
  - K711\_00xxx\_DELETE.ITM and R711\_00xxx\_DELETE.ITM files remove the ABAP code. The DELETE transport does not need to be imported, unless you stop the use of product entirely and want to remove the transports from the SAP systems. See [“Deleting the ABAP transport from the SAP system” on page 468](#).
3. Copy your transport files to the SAP Transport System data directory as follows:

**Remember:** You must not change the transport file name

Unicode transport

  - a. Copy the K711\_00xxxU.ITM file in the *cofiles* directory.
  - b. Copy the R711\_00xxxU.ITM file in the *data* directory.
4. To install the single IBM Tivoli Monitoring transport file on the SAP system, select one of the following file import options:
  - For the SAP system that is a Solution Manager 7.1 Service Pack 6 level, or later and is MAI configured, you must create the *transport.prop* file in the *usr/sap/SID/DVEBMGSinstancenumber/work* work directory of the SAP system. If the SAP system is a distributed system with ABAP SAP Central Services (ASCS), create the *transport.prop* file in the Central Instance (CI) *usr/sap/SID* directory. Then, add **MAI\_CONFIGURED = YES** entry in that file. This entry creates a **MAI\_CONFIGURED = YES** entry in the */IBMMON/ITM\_CNFG* table. You can now import the single IBM Tivoli Monitoring transport file on the SAP system.

**Note:** Before you import the single transport file, you must create the *transport.prop* file in the *usr/sap/SID/DVEBMGSinstancenumber/work* work directory of the SAP system and add **MAI\_CONFIGURED = YES** entry in that file. You must not edit the entry in the */IBMMON/ITM\_CNFG* table.

  - For all other SAP systems with basis version equal to 7.0, or later and Solution Manager V7.1 without MAI configuration, you must directly import the single IBM Tivoli Monitoring transport file.
5. Run the following command to import SAP transport:

```
tp addtobuffer ITMK711_00xxxU SID
pf=\usr\sap\sap\trans\bin\PROFILE_NAME
```

Where:

**SID**

Target SAP system ID.

**PROFILE\_NAME**

Name of the tp profile file. Make sure that the current tp parameter file is specified when you import the agent transport files from the command line. The tp parameter file is typically named TP\_DOMAIN\_SID.PFL. This file name is case-sensitive on UNIX systems.

**nnn**

Number for the target client where the agent runs and for which the user ID, IBMMON\_AGENT, authorization profile, and /IBMMON/AUTH, are defined.

Alternately, you can use the SAP STMS transaction to import the ITMK711\_00xxxU.ITM transport requests. Ensure that the following options are selected in the **Import Options** tab of the **Import Transport Request** window.

- **Leave Transport Request in Queue for Later Import**
- **Import Transport Request Again**
- **Overwrite Originals**
- **Overwrite Objects in Unconfirmed Repairs**

For the SAP Basis version, if the **Ignore Invalid Component Version** option is enabled, ensure that it is selected.

**Results**

Depending on your SAP release level, when you run the **tp import** command, you might receive return code 4, which does not indicate a problem. Receiving return code 4 is an expected result from the **import** command.

**Users and authorizations required by the SAP agent**

To safeguard against unauthorized access to the SAP system, you can assign authorizations to a user who logs in to the SAP system. These authorizations define the access levels for a user in the SAP system.

After you import the ABAP transport, the SAP agent creates the default user ID as IBMMON\_AGENT in the SAP system with the default password as ITMMYSAP. This user is a system user and the /IBMMON/AUTH authorization profile is associated with the user. The /IBMMON/AUTH profile and the IBMMON\_AGENT user are created after ABAP transport is imported. With the /IBMMON/AUTH profile, the IBMMON\_AGENT user can access transactions that are required to read performance data from the SAP system. Some examples of transactions that are used are as follows:

- CCMS alerts and administration
- Authorization for PI/XI message monitoring
- Solution Manager authorizations

You can create any other system type user for the agent. The user must be assigned the /IBMMON/AUTH profile.

To view and access data of SAP components, ensure that the user that is created for the agent has all the authorizations that are specified in the following table:

Table 59. The list of authorizations

| Components                                                                                                                                                  | Authorization objects                                      | Authorization description                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| General system authorizations that include the following components: <ul style="list-style-type: none"> <li>• SAP Instance</li> <li>• SAP System</li> </ul> | S_ADMI_FCD                                                 | To access the SAP system                                                                                                                                                                                                                                                                                         |
|                                                                                                                                                             | S_BDS_DS -BC-SRV-KPR-BDS                                   | To access the document set                                                                                                                                                                                                                                                                                       |
|                                                                                                                                                             | S_BTCH_JOB                                                 | To run operations on the background jobs                                                                                                                                                                                                                                                                         |
|                                                                                                                                                             | S_CCM_RECV                                                 | To transfer the central system repository data                                                                                                                                                                                                                                                                   |
|                                                                                                                                                             | S_C_FUNCT                                                  | To make C kernel function calls in the ABAP programs                                                                                                                                                                                                                                                             |
|                                                                                                                                                             | S_DATASET                                                  | To access files                                                                                                                                                                                                                                                                                                  |
|                                                                                                                                                             | S_RFC                                                      | To check RFC access. The S_RFC authorization object contains the following two subauthorizations: <ul style="list-style-type: none"> <li>• RFC1: To provide the authorizations for the RFC1 function group.</li> <li>• SDIFRUNTIME: To provide the authorizations for the SDIFRUNTIME function group.</li> </ul> |
|                                                                                                                                                             | S_RFACL                                                    | To check authorization for RFC users                                                                                                                                                                                                                                                                             |
|                                                                                                                                                             | S_RZL_ADM                                                  | To access Computing Center Management System (CCMS) for R/3 System administration                                                                                                                                                                                                                                |
|                                                                                                                                                             | S_TCODE                                                    | To check authorizations for starting the transactions that are defined for an application                                                                                                                                                                                                                        |
| S_TOOLS_EX                                                                                                                                                  | To display external statistics records in monitoring tools |                                                                                                                                                                                                                                                                                                                  |
| Authorizations for PI that include the SAP Process Integration                                                                                              | S_XMB_MONI                                                 | To access XI message monitoring                                                                                                                                                                                                                                                                                  |

Table 59. The list of authorizations (continued)

| Components                                                   | Authorization objects | Authorization description                                                              |
|--------------------------------------------------------------|-----------------------|----------------------------------------------------------------------------------------|
| Authorizations for MAI that include the SAP Solution Manager | AI_DIAGE2E            | To restrict E2E Diagnostics functions                                                  |
|                                                              | AI_LMDB_OB            | To access Landscape Management Database (LMDB) objects                                 |
|                                                              | SM_MOAL_TC            | To control the access to the alerting and monitoring functions in SAP Solution Manager |
|                                                              | SM_WC_VIEW            | To restrict access to specific UI elements in work centers of the Solution Manager     |
|                                                              | S_RFC_ADM             | To control rights for administering RFC destinations                                   |
|                                                              | S_RS_AUTH             | To specify analysis authorizations within a role                                       |
|                                                              | SM_APPTYPE            | To access Solution Manager app type                                                    |
|                                                              | SM_APP_ID             | To access applications provided in work centers                                        |

### Deleting the ABAP transport from the SAP system

If you choose to remove the SAP agent from your system, you must import delete transport to the SAP system. Delete transport deletes the SAP agent dictionary objects and function modules.

### Before you begin

Stop the SAP agent instance that is configured to monitor the SAP system.

If the SAP system is version 7.20 or later, you must add the following transport profile parameter: **tdirdelentions=true**. This transport profile parameter is available in tp version 375.57.68 and also in the R3trans version 6.14 release 700 or higher. For more information about removing transport requests from the SAP system, see [Deleting transport requests](#).

### Procedure

- Go to the following path:
  - On the Windows systems, *install\_dir*\TMAITM6\_x64\ABAP
  - On the Linux and AIX systems, *install\_dir*/*intrp*/sa/ABAP, here *intrp* must be **1x8266** or **aix526**.
- Copy the transport files into the SAP environment.
- Copy the K711\_00xxx\_DELETE and R711\_00xxx\_DELETE files to the SAP Transport System data directory as follows:
  - Copy the K711\_00xxx\_DELETE file to the cofiles directory.
  - Copy the R711\_00xxx\_DELETE file to the data directory.
- Run the following commands to import the delete transport:
  - tp addtobuffer ITMK711\_00xxx\_DELETE SID pf=\usr\sap\trans\bin\PROFILE\_NAME**
  - tp import ITMK711\_00xxx\_DELETE SID client=nnn U16 pf=\usr\sap\trans\bin\PROFILE\_NAME**



Where:

**SID**

Target SAP system ID.

**PROFILE\_NAME**

Name of the tp profile file.

**nnn**

Number for the target client where the agent is to run.

## Verifying agent configuration

After you install the SAP agent, you must verify the agent configuration by downloading, copying, and verifying the NetWeaver RFC SDK V7.20 library. You must also verify the configuration of Solution Manager V7.1 with MAI\_Monitoring, verify MAI Alerts, and verify the configuration setting specific to third-party component.

Verify the agent configuration by completing the following procedures:

- [“Downloading the NetWeaver RFC SDK V7.20 library” on page 469](#)
- [“Copying the NetWeaver RFC SDK V7.20 library in SAP agent setup” on page 470](#)
- [“Verifying the NetWeaver RFC SDK V7.20 library” on page 470](#)
- [“Verifying the configuration of Solution Manager V7.1 with MAI-Monitoring” on page 471](#)
- [“Verifying MAI Alerts” on page 472](#)
- [“Verifying configuration settings specific to third-party component” on page 472](#)

### Downloading the NetWeaver RFC SDK V7.20 library

Download the NetWeaver RFC SDK V7.20 library after you finish installing the SAP agent. All the files that are related to the NetWeaver RFC SDK V7.20 library are available for download on the SAP website.

### Procedure

Follow this procedure to download the Net Weaver RFC SDK V7.20 library.

1. Log in to SAP Marketplace by using the following URL:  
<http://service.sap.com>
2. Click **SAP Support Portal**.
3. Enter your Service Marketplace user name and password.
4. Click **Software Downloads** and expand the **Support Packages and Patches** link.
5. Click **Browse our Download Catalog**, and then click **Additional Components**.
6. Click **SAP NetWeaver RFC SDK**, and then click **SAP NetWeaver RFC SDK 7.20**.
7. Select the operating system where you have the SAP agent.
8. Download the \*.SAR file on your computer.
9. To extract the SAP Netweaver RFC SDK \*.SAR file by using the SAPCAR utility that is provided by SAP, run the following command:

```
sapcar -xvf SAP NetWeaver RFC SDK File Name.SAR
```

**Note:** You can download the SAPCAR utility from the SAP website.

10. Go to the lib folder inside the extracted folder.

### What to do next

Copy the NetWeaver RFC SDK V7.20 library in to the SAP agent setup.

## Copying the NetWeaver RFC SDK V7.20 library in SAP agent setup

The NetWeaver RFC SDK V7.20 library contains files that you must manually copy in the SAP agent setup location.

### Procedure

1. Go to the directory where you downloaded the NetWeaver RFC SDK V7.20 library.
2. Copy the files to the SAP agent setup location.

- For Windows 64-bit operating systems you must copy the following files:
  - icuin34.dll
  - libicudecnumber.dll
  - libsapucum.dll
  - icudt34.dll
  - icuuc34.dll
  - sapnrwrfc.dll

You must copy the files to *install\_dir*\TMAITM6\_x64 location.

- For operating systems other than Windows, you must copy the files to the *install\_dir*/*intrap*/sa/lib location, where *intrap* is the operating system code (aix526, li6263, sol606). You must copy the following files:
  - libsapnrwrfc.so
  - libicudecnumber.so
  - ibicuuc34.a
  - libicui18n34.a
  - libicudata34.a
  - libsapucum.so

### What to do next

Verify the version of the NetWeaver RFC SDK V7.20 library that is downloaded.

### Verifying the NetWeaver RFC SDK V7.20 library

You must verify the version of the file after you copy the extracted file.

### Procedure

- **Windows** To verify the version of the file, complete the following steps:
  - a) Right-click `sapnrwrfc.dll` and click **Properties**.
  - b) Click the **Version** tab.
  - c) In the **Product Version** section, ensure that you have the following version: 720, patch 514, changelist 1448293, or later.
- **Linux** **UNIX** To verify the version of the file, complete the following steps:
  - a) Go to the `lib` folder in the extracted \*.SAR file.
  - b) Run the following command: **strings libsapnrwrfc.so | grep SAPFileVersion**
  - c) You must see the following message: `[root@IBMSAP2V6 lib]# strings libsapnrwrfc.so | grep SAPFileVersion GetSAPFileVersion #[%]SAPFileVersion: 7200, 514, 22, 6206 .GetSAPFileVersion`

**Note:** The message shows that this library has the version 720 patch 514, or later.

## Verifying the configuration of Solution Manager V7.1 with MAI-Monitoring

To receive data for MAI Alerts, you must verify whether the Solution Manager V7.1 is configured correctly.

### About this task

You can use Solution Manager V7.1 with MAI-Monitoring and Alerting Infrastructure to monitor the Managed Systems. Solution Manager V7.1 monitors itself and the satellite systems. Each satellite system has a plug-in and diagnostics agents. Diagnostics agents fetch the data for Host or Operating System level. Each host can have multiple diagnostics agents for different Solution Managers monitoring the host. Following are the keywords that are used in Solution Manager MAI Monitoring:

- Metrics: Data from the satellite systems.
- Alerts: Notifications that are based on some crossovers of threshold values that can be configured.
- Incident: Alerts that are converted into tickets and assigned to any user.

To verify the configuration of Solution Manager V7.1 with MAI monitoring, you must verify the basic settings, global level settings, and template level settings.

### Procedure

1. To verify the basic settings, enter the Transaction Code: SOLMAN\_SETUP and click **Enter**.  
Ensure that all the LEDs are green in the following tabs:
  - Overview
  - Basic Configuration
  - Managed System Configuration

**Note:** There are different categories of Managed Systems such as Technical Systems, Technical Scenarios, Host, Database, Instance, PI Domain, Technical Component, and Connection. You must configure these Managed Systems according to business requirements. The MAI Alerts are based on the Managed Systems that you configured.
2. Enter the Transaction code: SE38 and click **Enter**.
3. Provide the program name as RTCCTOOL and run the report.  
Ensure that all the LEDs are green in the output.
4. To verify the global level settings, enter the Transaction code: SOLMAN\_WORKCENTER and click **Enter**.  
Ensure that all the LEDs are green in the following tabs:
  - Overview
  - Configure Infrastructure
  - Pre-requisites
  - Configure
5. Verify whether the **Global Settings** for **Notification** status is **Active**.
6. To verify the template level settings, enter the Transaction Code: SOLMAN\_SETUP and click **Enter**.  
In **Technical Settings**, in the **Auto-Notifications** list, ensure **Active** is selected.

**Note:** For initial troubleshooting, ensure that email notifications are active.
7. For MAI system monitoring, verify the configuration of End-User Experience Monitoring (EEM) by using the following steps:
  - a) Enter the Transaction code: SE37 and press **Enter**.
  - b) Enter **AI\_EEM\_LIST\_ALL\_SCENARIOS** in the **Function Module name** field and press F8.  
There must be an entry for End-User Experience Monitoring (EEM).

## Verifying MAI Alerts

To ensure that Solution Manager MAI is configured correctly for monitoring the MAI Alert Inbox in Technical Monitoring, you must verify that you receive MAI Alerts as output.

### Procedure

1. Enter the Transaction code SOLMAN\_WORKCENTER and click **Enter**. Check whether you can view MAI Alerts in the Solution Manager MAI Alert Inbox under Technical Monitoring.
2. Check for BAdi implementation by using the following steps:
  - a) Enter the Transaction code: SE19 and click **Enter**.
  - b) Enter /IBMMON/ITM\_IMPL\_ALRTINBX in the **Enhancement Implementation** field.
  - c) Click **Display** and check if BAdi implementation is active in **Runtime Behavior** section.
3. Check whether the database /IBMMON/ITM\_ALIX contains MAI Alerts by using the following steps:
  - a) Enter the Transaction code: SE16 and press **Enter**.
  - b) In the **Table Name** field, enter /IBMMON/ITM\_ALIX and run it. Ensure that you are receiving MAI Alerts in the table.
4. Enter the Transaction code: SE37 and click **Enter**.
5. In the **Function Module Name** field, enter /IBMMON/ITM\_MAIALRT\_INX and press F8.  
You must see MAI Alerts as output.

### What to do next

If you are not able to view MAI Alerts in the /IBMMON/ITM\_ALIX database, you must verify the settings in the Third-Party Component.

### Verifying configuration settings specific to third-party component

If you are not able to view MAI Alerts, then you must verify the settings in the third-party component.

### Procedure

1. Verify that Third-Party Component is active.
2. Verify that in **OS Adapter**, under **BAdi Implementation**, **Alert Reaction** is available. If **Alert Reaction** is not available, remove the default settings, and select the **BAdi implementation - Alert Reaction**.
3. Check the template settings by using the following steps:
  - a) Verify the settings that are used to transfer specific alerts to the Third-Party System such as SAP ABAP 7.0.0.
  - b) Select **Expert Mode**, select **Alerts**, and then click **Third Party Component**.  
Ensure that you are able to view the Alert Reaction BAdi name.  
**Note:** Ensure that the latest SAP notes are implemented. For Solution Manager V7.1 Service Pack 8, check if the following notes are implemented:
    - <https://service.sap.com/sap/support/notes/1959978>
    - <https://service.sap.com/sap/support/notes/1820727>
4. If you are not able to view MAI Alerts in the /IBMMON/ITM\_MAIALRT\_INX database, you must run the following Solution Manager MAI configurations steps for Third-Party Component:
  - a) Enter the Transaction code: SOLMAN\_SETUP and click **Enter**.
  - b) In **Technical Monitoring**, select **System Monitoring**.
  - c) Click **Configure Infrastructure** tab and then click **Default Settings** tab.
  - d) Click **Third Party Components** tab and then click **Edit**.
  - e) Select **Active** from the list.
  - f) Ensure that scope filter is set as **All alerts, Events and Metrics (with Internal Events)** for the selected connector.

**Note:** OS Command Adapter is also one of the methods to push data to the third-party connector. To configure the OS Command Adapter, read the configuration detail settings in the How-to guide for OS Command Adapter.

## Advanced installation and configuration of the SAP agent

The following advance installation and configurations are specific to the SAP agent.

- [“SAP function module” on page 473](#)
- [“SAP user IDs” on page 474](#)
- Utilities for the SAP agent
- [“SAP RFC connections” on page 474](#)
- [“Test Connection feature” on page 484](#)
- [“Optional advanced configuration in SAP” on page 476](#)
- [“CEN CCMS reporting” on page 482](#)
- [“Uninstalling the Advanced Business Application Programming \(ABAP\) transport from the SAP system” on page 483](#)

**Note:** The advance installation and configuration of the SAP agent contains references to IBM Tivoli Monitoring to make the documentation compatible with ABAP transport custom transaction code UI.

### SAP function module

When the data volume is high on the SAP server, you might experience problems with certain widgets to cause a slow response time from the server. If the widgets are not critical, you can disable the associated SAP function module.

By default, the SAP agent function modules are enabled. However, the following function modules are unavailable by default:

- HTTP services under the SYS subnode (/IBMMON/ITM\_HTTP\_SRVS)
- XML messages under the PI/XI subnode (/IBMMON/ITM\_SXMB\_MONI\_NEW)
- Sync/Async communication under the PI/XI subnode (/IBMMON/ITM\_SYN\_ASYNC\_COMM)
- qRFC inbound queue details under the Sys subnode (/IBMMON/ITM\_QIN\_QDETAILS)

After disabling the SAP function module, if you select a widget, data isn't displayed on IBM Cloud App Management UI. Therefore, you avoid any performance-related problems.

### Enabling the SAP agent function module

If you have previously disabled the SAP agent function module to resolve performance problems, then you can enable the function module too.

#### Procedure

1. Log on to the SAP system.
2. Run the SE16 transaction code.
3. Enter table name as /IBMMON/ITM\_CNFG.
4. Select the row to delete and press **shift + F2** to delete the entry.
5. Click **Save**.

### Disabling the SAP function module

Some widgets might cause a slow response from the SAP server so you can disable the SAP function module to improve the server performance.

#### Procedure

1. Log on to the SAP system.
2. Run the SE16 transaction code.

3. Enter table name as /IBMMON/ITM\_CNFG.
4. Press **F5** to create a new entry.
5. Enter the name of the SAP function module in the **PARM NAME** field.
6. Enter No in the **VALUE CHAR** field.
7. Click **Save**.

### **SAP user IDs**

This section provides information about SAP user IDs and permissions that are required by the SAP agent.

User IDs support the following purposes:

- [“SAP RFC connections” on page 474](#)
- [“Basic agent monitoring” on page 474](#)

### **SAP RFC connections**

The SAP agent uses Remote Function Calls (RFC) connections for internal Centralized Computing Center Management (CCMS) polling and CCMS alert data collection. This behavior is specific to the SAP RFC architecture.

The SAP agent opens one dedicated RFC connection to the SAP system that is monitored by the agent. The SAP system then opens one internal connection per application server for data collection through function modules and programs. If CCMS alerts are collected by the agent, the SAP system opens one more (system internal) RFC connection to each application server for this collection thread. When data collection starts, one RFC connection for the agent is opened. Then, up to twice the number of SAP application servers for more internal system RFC connections are opened.

You must ensure that the instance that is monitoring can accommodate the additional RFC sessions, especially in large systems with 10, or more instances. When the anticipated RFC load for monitoring might adversely affect system performance and tolerances, adjust the SAP profile parameter. Contact your SAP Administrator and see the following SAP Notes:

- Terminal Sessions (default setting: 200) 22099
- Communication/Gateway/Conversation Settings 887909 316877 384971

### **Basic agent monitoring**

The SAP agent creates an IBMMON\_AGENT in the SAP system when the agent transport is imported.

This user ID is IBMMON\_AGENT with the default password ITMMYSAP. It is preconfigured to be Communication Type user-only and to use the /IBMMON/AUTH authorization profile. This profile, which is created at transport import time, contains the minimal set of permissions to run the agent Advanced Business Application Programming (ABAP) code. Also, this profile accepts a set of limited actions on your SAP system.

If this user ID name is unacceptable, for example, if it violates your naming conventions that are used during installation, you can create a different user ID. The user ID can be any allowable SAP user ID, but it requires the complete set of permissions in the /IBMMON/AUTH profile. The user ID requires Communication Type user-only access.

The default user ID provides sufficient authority only for the following purposes:

- Monitoring and data collection
- Closing Computing Center Management System (CCMS) alerts
- Enabling, disabling, and resetting gateway statistics
- Resetting Oracle database statistics

If you choose to limit the action capabilities of the agent, you can remove some of the action permissions such as closing CCMS alerts.

To access data on the IBM Cloud App Management UI Portal for specific components, ensure that you have appropriate authorizations. Following table lists the authorizations that are required to access the data from different sub nodes:

Table 60. The list of authorizations

| Sub nodes                                                                                                                                        | Authorization objects               | Authorization description                                                                                                                                                                                                                                                                                        |
|--------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| General system authorizations that include the following sub nodes: <ul style="list-style-type: none"> <li>• Ins</li> <li>• Sys</li> </ul>       | S_ADMI_FCD                          | To access the System                                                                                                                                                                                                                                                                                             |
|                                                                                                                                                  | S_BDS_DS -BC-SRV-KPR-BDS            | To access the Document Set                                                                                                                                                                                                                                                                                       |
|                                                                                                                                                  | S_BTCH_JOB                          | To run operations on the background jobs                                                                                                                                                                                                                                                                         |
|                                                                                                                                                  | S_CCM_RECV                          | For transferring the Central System Repository data                                                                                                                                                                                                                                                              |
|                                                                                                                                                  | S_C_FUNCT                           | To make C calls in the ABAP programs                                                                                                                                                                                                                                                                             |
|                                                                                                                                                  | S_DATASET                           | To access files                                                                                                                                                                                                                                                                                                  |
|                                                                                                                                                  | S_RFC                               | To check RFC access. The S_RFC authorization object contains the following two subauthorizations: <ul style="list-style-type: none"> <li>• RFC1: To provide the authorizations for the RFC1 function group.</li> <li>• SDIFRUNTIME: To provide the authorizations for the SDIFRUNTIME function group.</li> </ul> |
|                                                                                                                                                  | S_RFACL                             | For RFC User                                                                                                                                                                                                                                                                                                     |
|                                                                                                                                                  | S_RZL_ADM                           | To access Computing Center Management System (CCMS): System Administration                                                                                                                                                                                                                                       |
|                                                                                                                                                  | S_TCODE                             | To check Transaction Code at Transaction Start                                                                                                                                                                                                                                                                   |
| S_TOOLS_EX                                                                                                                                       | To access Tools Performance Monitor |                                                                                                                                                                                                                                                                                                                  |
| Authorizations for Solution manager that include the following sub nodes: <ul style="list-style-type: none"> <li>• Lds</li> <li>• Sol</li> </ul> | D_MD_DATA -DMD                      | To view Data Contents of Master Data                                                                                                                                                                                                                                                                             |
|                                                                                                                                                  | D_SOLMANBU                          | To access a Session Type of the Solution Manager                                                                                                                                                                                                                                                                 |
|                                                                                                                                                  | D_SOLM_ACT                          | To access a Solution in the Solution Manager                                                                                                                                                                                                                                                                     |
|                                                                                                                                                  | D_SOL_VSBL                          | To view a Solution in the Solution Manager                                                                                                                                                                                                                                                                       |
|                                                                                                                                                  | S_CTS_SADM                          | To view System-Specific Administration (Transport)                                                                                                                                                                                                                                                               |
|                                                                                                                                                  | S_TABU_RFC                          | To view Client Comparison and Copy: Data Export with RFC                                                                                                                                                                                                                                                         |
| Authorizations for PI that includes the PI sub node                                                                                              | S_XMB_MONI                          | To access XI Message Monitoring                                                                                                                                                                                                                                                                                  |

Table 60. The list of authorizations (continued)

| Sub nodes                                             | Authorization objects | Authorization description                              |
|-------------------------------------------------------|-----------------------|--------------------------------------------------------|
| Authorizations for MAI that includes the SIm sub node | AI_DIAGE2E            | To access Solution Diagnostics end-to-end analysis     |
|                                                       | AI_LMDB_OB            | To access Landscape Management Database (LMDB) Objects |
|                                                       | SM_MOAL_TC            | To access Monitoring and Alerting                      |
|                                                       | SM_WC_VIEW            | To access Work Center User Interface Elements          |
|                                                       | S_RFC_ADM             | To access Administration options for RFC Destination   |
|                                                       | S_RS_AUTH             | To access BI Analysis in Role                          |
|                                                       | SM_APPTYPE            | To access Solution Manager App Type                    |
|                                                       | SM_APP_ID             | To access applications provided in Work center         |

### Central User Administration (CUA)

The Central User Administration (CUA) is used to monitor a SAP system.

### Procedure

To use the predefined user ID and authorization role to monitor a SAP system set-up with Central User Administration, complete one of the following steps:

- Install the transport into the Central User Administration parent logical system client.
- Manually create the user ID or role in the client where you want to install the transport. The user ID or role is in the client where the transport is installed (imported).
- Manually create the user ID or role in the Central User Administration parent logical system client. Then, distribute the user ID or role to the client where the agent runs.
- Manually create the user ID or role in the Central User Administration parent logical system client and run the agent in this client.

### Optional advanced configuration in SAP

You can configure the SAP agent by using standard SAP or agent-provided SAP functions.

Use agent-provided transactions in SAP to customize a number of agent behaviors. After you run the /n/ IBMMON/ITM\_CONFIG transaction to access the main configuration menu in SAP, select one of the following configuration options:

- [“Copy, back up, restore feature and transactions” on page 477](#)
- [“Copy, back up, and restore data by using transactions” on page 477](#)
- [“Command-line utility tool” on page 478](#)
- [“Running the command-line utility on a Windows environment” on page 479](#)
- [“Running the command-line utility on a Non-Windows environment” on page 479](#)
- [“Alerts maintenance” on page 480](#)
- [“Select monitor sets and monitors transaction” on page 480](#)
- [“Configure Dialog Step Response Threshold in the SAP system” on page 481](#)



**Note:** You must preface all /IBMMON/ITM\* transactions with /n.

Configuration changes made in these transactions are used immediately by the SAP agent except for those changes that are made to maintain managed groups. When the managed group configuration changes, the changes are discovered by the SAP agent at the next heartbeat.

Use SAP standard functions to complete the following configuration: [“Configure Dialog Step Response Threshold in the SAP system”](#) on page 481

### ***Copy, back up, restore feature and transactions***

The copy, back up, and restore features are available to you after you log on to the SAP server and run the following transaction: /n/IBMMON/ITM\_CONFIG.

Copy, backup, and restore operations allow you to copy, backup, and restore the IBM Tivoli Monitoring configuration data.

Use this feature to select from the following functions and to save the IBM Tivoli Monitoring configuration data:

- **Copy**

Use this feature to copy the IBM Tivoli Monitoring configuration settings from one SAP server to another SAP server. For example, you might want to copy the IBM Tivoli Monitoring configuration settings from agent **a1** to SAP server instance SAP2. This agent runs on system **m1** and is configured for SAP server instance SAP 1. All the IBM Tivoli Monitoring configuration settings, except the SAP server instance monitoring settings are copied to the target SAP system. You implement the copy feature by using either the command line utility or the SAP GUI.

- **Backup**

You can store agent-specific configurations that you completed on the SAP server by taking a backup of the system. Use this feature to save IBM Tivoli Monitoring specific configuration settings on the SAP system. You use the /IBMMON/ITM\_CONFIG transaction to enter the settings. The backup file is stored in the work directory on the SAP server to the following path: /usr/sap//DVEBMGS/work.

- **Restore**

Use this feature to restore IBM Tivoli Monitoring configuration data on the SAP server from the work directory. You can restore the IBM Tivoli Monitoring configuration data on the same SAP server where you completed the backup procedure of this configuration data or another SAP server. You can restore IBM Tivoli Monitoring configuration data to specific SAP and IBM Tivoli Monitoring tables. Configuration files are stored with a date and time stamp so you can select the point to which you want to restore your files.

Agent-specific configurations include configuration settings in the /IBMMON/ITM\_CONFIG transaction in SAP. You can complete the following configuration procedures:

- Sample the frequency for alerts
- Enable specific alerts
- Store log file names
- Manage group definitions
- Select monitor sets and monitors
- Select SAP instances for monitoring purposes

### ***Copy, back up, and restore data by using transactions***

On the SAP user interface, you can copy, back up, and restore data by using the /n/IBMMON/ITM\_CONFIG transaction.

### **Before you begin**

Use the copy, backup, and restore procedures to copy the IBM Tivoli Monitoring configuration settings from one SAP server to another SAP server. All the IBM Tivoli Monitoring configuration settings, except the SAP server instance monitoring settings are copied to the target SAP system.

## Procedure

To copy, back up, and restore your data on SAP complete the following steps:

- To copy your data on SAP complete the following steps:
  - a. Enter the target SAP system ID and the existing file name as source system id\_\_<filename>date\_time. The /IBMMON/ITM\_COPY transaction creates an IBM Tivoli Monitoring configuration file in the work directory with the filename as SAP target SAP system id\_\_<filename>\_date\_time.
  - b. Click **Execute** to copy the IBM Tivoli Monitoring configuration data to the file.
  - c. Click **Back** or **Cancel** for returning to the previous IBM Tivoli Monitoring configuration screen. The expected input parameters are **Target System id** and **filename** which are to be copied.
- To back up your data on SAP complete the following steps:
  - a. Log on to the SAP server and start the /IBMMON/ITM\_CONFIG transaction.
  - b. Select **Backup** and enter the backup filename.
  - c. Enter the backup filename.

The file name is stored as sys\_id\_<filename>\_date\_time.
  - d. Click **Execute** to run the backup and to store the file on the Application Server.

**Note:** The backup file is stored in the work directory of the application server.
  - e. Click **Back** or **Cancel** for returning to the previous IBM Tivoli Monitoring configuration screen.
- To restore your data on SAP complete the following steps:
  - a. Log on to the SAP server and start the /IBMMON/ITM\_CONFIG transaction.
  - b. Select **Restore**.
  - c. Enter the filename to restore as sys\_id\_<filename>\_date\_time.
  - d. Click **Execute** to restore IBM Tivoli Monitoring configuration data.
  - e. Click **Back** or **Cancel** for returning to the previous IBM Tivoli Monitoring configuration screen.

### Command-line utility tool

You can use the command-line utility tool to copy, backup, and restore IBM Tivoli Monitoring configuration data on the SAP server.

You can run the command-line utility tool on Windows and Non-Windows environment. See [“Running the command-line utility on a Windows environment” on page 479](#) and [“Running the command-line utility on a Non-Windows environment” on page 479](#).

#### • Copy

Run the **backup** command to copy the IBM Tivoli Monitoring configuration file from the agent directory SAP server instance sap1 to sap2. Enter the file name and sap1 as the source system from the sap1 agent directory. Then, the ABAP function is called that copies the IBM Tivoli Monitoring settings from this file to the IBM Tivoli Monitoring configuration file for Sap2. Now select **Copy** from the sap1 agent directory utility tool and enter a file name and sap2 as the target SAP system.

#### • Backup

After running the command-line utility tool, select the **Backup** option. Then, you need to enter the file name and the SAP system ID. The tool calls the /IBMMON/ITM\_BACKUP SAP function module. The function module reads the specific IBM Tivoli Monitoring configuration settings that are stored in tables and stores them with a row and column separator. Then, the command-line utility tool reads the string and writes the data into a file. The file name that is generated has the following format: ID>\_<filename>><date&time>. This file is stored in the directory where the utility program is stored.

#### • Restore

After you run the command-line utility tool, enter the file name to restore and the target SAP system where you want to restore the file. The command-line utility tool reads the file from the agent directory

and calls the /IBMMON/ITM\_RESTORE SAP function module. Then, the tool passes the IBM Tivoli Monitoring configurations as a string. The SAP function module updates the specific IBM Tivoli Monitoring tables and restores the specific IBM Tivoli Monitoring configurations.

### *Running the command-line utility on a Windows environment*

You can run the command-line utility on a Windows environment to complete copy, backup, and restore procedures.

#### **Procedure**

1. Depending on your operating system, complete one of the following procedures:
  - For a 64-bit operating system, set the CANDLEHOME path by using command **set CANDLE\_HOME = C:\IBM\APM** and run the **ksacopybackuprestore.bat** command from the following path:  
%candle\_home%\ TMAITM6x64.
2. To create a backup file, complete the following steps:
  - a) Select **Backup** and enter the file name and source SAP system name.
  - b) The backup file is created with the following format: SYS ID>\_<filename>\_<date&time>.
3. To restore the file, complete the following steps:
  - a) Select **Restore** and enter the target SAP system name.
  - b) Enter the file name.
4. To copy the file, complete the following steps:
  - a) From the source agent, select **Backup** and create a backup file.
  - b) Copy the backup file from the source agent directory to the target agent directory.
  - c) From the source directory, run the command-line utility tool and select **Copy**.
  - d) Enter the file name and the target SAP system.

### *Running the command-line utility on a Non-Windows environment*

You can run the command-line utility on a Non-Windows environment to complete copy, backup, and restore procedures.

#### **Procedure**

1. Run the following command from /candle\_home/<arch>/sa/shell path.

```
ksacopybackuprestore.sh
```

2. To create a backup file, complete the following steps:
  - a) Select **Backup** and enter the file name and source SAP system name.
  - b) The backup file is created with the following format: SYS ID>\_<filename>\_<date&time>.  
The backup file is saved to this location: %candlehome% / arch /sa/bin.
3. To restore the file, complete the following steps:
  - a) Select **Restore** and enter the target SAP system name.
  - b) Enter the file name.
4. To copy the file, complete the following steps:
  - a) From the source agent, select **Backup** and create a backup file.
  - b) Copy the backup file from the source agent directory to the target agent directory.
  - c) From the source directory, run the command-line utility tool and select **Copy**.
  - d) Enter the file name and the target SAP system.

### ***Alerts maintenance***

You can modify alerts that are generated by IBM Tivoli Monitoring by changing their status and thresholds.

This transaction is used to enable or disable alerts that are generated by IBM Tivoli Monitoring and to set warning and critical thresholds. All alerts that are generated by IBM Tivoli Monitoring are shown with their status and threshold values.

When you modify alert status and thresholds, the modified values are used at the next sample time.

### **Default sample period maintenance**

The default sample period provides information about real-time reporting for certain attribute groups.

Some attribute groups have an implicit date and time for each record in the group. For example, the R/3\_Abap\_Dumps attribute group reports the create time for the dump and the R/3\_System\_Log attribute group reports the create time for the log entry. These records have a date and time field. You can obtain a report for a short history of the table instead of just the most recent information. This time interval is the time span for data collection and is used as the real-time interval when the data is collected. The /IBMMON/ITM\_PERIOD transaction defines a default sample period (time span for real-time reporting) for each of these attribute groups. The sample period identifies the length of the data sample period that starts from the current time and works back in time.

### **Log file name maintenance**

Specific log files that are matched only to instances are included in IBM Tivoli Monitoring reports with log file information.

This transaction is used to identify which log files to consider for inclusion in IBM Tivoli Monitoring reports that contain log file information. All log files with a name that matches the specified name patterns on the specified instances are included in the report at the next data collection interval.

### **Managed group maintenance**

The Managed Group names transaction monitors and processes specific transactions in the SAP system.

Use this transaction to maintain IBM Tivoli Monitoring Managed Group definitions. All Managed Group names are passed to the IBM Cloud App Management UI Portal and shown in the Managed System Selection Lists. At the time of data collection, only data that matches the Attribute selection conditions is sent to the SAP agent. This data is shown in reports or used for evaluation in situations and policies.

You use Managed Groups to monitor subsets of information in the SAP system. You focus only on the parts of the SAP system in which you are interested and you ignore other parts that do not concern you. For example, if you are only interested in the response time of transactions that are part of the Financial Application, you create a Managed Group that is named Financials. Then, you include only Financial transaction codes in it. Whenever the Financial Managed Group is processed by the Tivoli Enterprise Portal only information that contains the specified transaction codes is considered when a report is shown, situation or policy is evaluated..

### ***Select monitor sets and monitors transaction***

Use the select monitor sets and monitors transaction to edit the Centralized Computing Central Management (CCMS) alerts configuration. For example, you can turn off CCMS alert collection completely.

This transaction is used to select the CCMS monitors from which IBM Tivoli Monitoring retrieves alerts. By default, the Entire System Monitor is selected the first time that this window is shown. You can change the monitor set, the monitor, or both the monitor set and monitor, and then save the configuration. You can select a maximum of three monitors for which to collect CCMS alerts.

To turn off CCMS alert collection completely, clear the check boxes for all of the monitors and save this configuration.

The agent that is already running reads this configuration and collects the CCMS alerts for the monitors that you selected. However, any CCMS alerts that were already collected by the agent before changing the CCMS alerts configuration remain with the agent and IBM Tivoli Monitoring.

In addition to selecting monitors and monitors sets, this transaction specifies the number of occurrences of an alert type to retrieve. Also, it helps you to decide whether to automatically close the older occurrences of the alerts that are not retrieved.

### Configure Dialog Step Response Threshold in the SAP system

You can configure a Dialog Step Response Threshold for any transaction by running the SE16 transaction.

#### Procedure

1. In the **Table Name** field, type /IBMMON/ITM\_TRSH, and then select **Table Contents (F7)** to access the table.
2. To view the current threshold settings, select **Execute (F8)**. The transaction names are shown under **WORKLOAD** column; the threshold values are shown under the **THRESHOLDWORKLOAD** column.
3. To add threshold setting, select **Create (F5)**. Type the transaction name in the field. The following wildcards are accepted for the **WORKLOAD** value:
  - \* matches multiple characters
  - + matches any single character
4. Type the threshold value, in milliseconds, in the **THRESHOLD** field. Select **Save** to save this setting. New and changed threshold values do not take effect immediately, but take effect under either of the following conditions:
  - The agent is restarted.
  - The agent reopens its RFC connection to the SAP system. This procedure occurs every 12 heartbeats, which, by default, is about every 2 hours and 10 minutes.

#### Results

The value that is entered for the **Threshold** column is returned in the Dialog Step Response Threshold attribute of the R/3\_Transaction\_Performance attribute group.

### Batch Job Operations

You can fetch all the Batch Jobs within a specified time interval.

#### Procedure

Follow the steps after [“Importing the ABAP transport on the SAP system”](#) on page 462.

**Remember:** Critical Constant is set for all the batch jobs.

1. To fetch all Active and Canceled Batch Jobs within a specified time interval. Add the following entry in /IBMMON/ITM\_CNFG table.

| Table 61. /IBMMON/ITM_CNFG |            |
|----------------------------|------------|
| PARAM_NAME                 | VALUE_CHAR |
| BATCH_JOBS_PERF            | YES        |

2. To fetch all Canceled jobs within a specified time interval and all Active jobs irrespective of time interval. Add the following entry in /IBMMON/ITM\_CNFG table.

| Table 62. /IBMMON/ITM_CNFG |              |
|----------------------------|--------------|
| PARAM_NAME                 | VALUE_CHAR   |
| BATCH_JOBS_PERF            | YES_LONG_RUN |

- To fetch all Batch Jobs within a specified time interval and all Active Batch Jobs irrespective of time interval.

Add the following entry in /IBMMON/ITM\_CNFG table.

| Table 63. /IBMMON/ITM_CNFG |            |
|----------------------------|------------|
| PARAM_NAME                 | VALUE_CHAR |
| BATCH_JOBS_PERF            | YES_ALL    |

**Note:**

- If the configuration parameter is not added, it fetches all Batch Jobs within a specified time interval without Critical Constant set.
- Number of rows that are fetched is always equal to value of Critical Constant set in Transaction Code /n/IBMMON/ITM\_CONFIG.

**Improving /IBMMON/ITM\_MAIALRT\_INX Function Module's performance**

You can enhance the /IBMMON/ITM\_MAIALRT\_INX Function Module's performance for SAP agent.

**Procedure**

Follow the steps to improve the /IBMMON/ITM\_MAIALRT\_INX function module's performance.

- Log on toSAP agent GUI.
- Run SE16 transaction code and enter the table name as /IBMMON/ITM\_CNFG and press F7.
- Press F5 or click **Create Entries** and add the following entry in the IBMMON/ITM\_CNFG table.

| Table 64. /IBMMON/ITM_CNFG |            |
|----------------------------|------------|
| PARAM_NAME                 | VALUE_CHAR |
| MAI_ALERTS_PERF            | YES        |

**Note:**

- If the Critical Constant is not set in the Transaction Code - /N/IBMMON/ITM\_CONFIG, then default value is 2500.
  - This process is only applicable for fetching the MAI Alerts from the SAP system where thePERIOD\_START and PERIOD\_END is initial.
- Remember:** Now the Function Module /IBMMON/ITM\_MAIALRT\_INX fetches the number of MAI Alerts equivalent to the Critical Constant set in the Transaction Code - /N/IBMMON/ITM\_CONFIG.
- If this entry in the /IBMMON/ITM\_CNFG is not created by default, then the 2500 latest MAI alerts are fetched.
  - The number of rows that are fetched is always equal to value of Critical Constant set in Transaction Code /n/IBMMON/ITM\_CONFIG.

**CEN CCMS reporting**

Centralized (CEN) Computing Center Management System (CCMS) is a SAP monitoring capability.

Use this capability to report CCMS alerts for multiple SAP systems to a central monitoring hub. You monitor the SAP environment from one CCMS console. Centralized CCMS reporting is best used in the following environments:

- Primarily a CCMS operation where CCMS alerts are the only monitoring data needed.
- Centralized CCMS is part of the SAP environment.
- Large SAP environments with many SAP systems such as ISV and ISP.
- IBM Tivoli Monitoring V5.x integration with SAP agent CCMS adapters.
- Collect alerts from non-ABAP SAP components and application servers.

The SAP agent supports Centralized CCMS for reporting alerts only. Then, you place one SAP agent on a Centralized SAP system and view CCMS alerts for the entire SAP environment. This support is provided in the following ways:

- When reporting CCMS alerts, the agent checks if the alerts are associated with the SAP system that is directly monitored by the agent. If the agent determines that an alert belongs to a different SAP system, it assumes Centralized CCMS and automatically creates more R3\_Group managed systems.
- The <local\_SID>-All\_CCMS\_alerts:Grp managed system is used to report the complete set of alerts from all remote SAP systems. The value of <local\_SID> is the system identifier for the SAP system that is directly monitored. For example, if the local SAP system is QA1, this group name would be QA1-All\_CCMS\_alerts:Grp.
- The <local\_SID>-<remote\_SID>\_CCMS\_alerts:Grp managed system is used to report all alerts for one remote SAP system. The value of <local\_SID> is the system identifier for the SAP system that is directly monitored. The value of <remote\_SID> is the system identifier for the remote SAP system. For example, if the local SAP system is QA1 and the remote SAP system is QA2, this group name would be QA1-QA2\_CCMS\_alerts:Grp.
- Each of these managed systems in the Navigator tree has the complete set of widgets under it, but only the Alerts widgets have meaningful data.

The SAP agent maintains its definitions of Centralized CCMS groups in the Advanced Business Application Programming (ABAP) code in the directly managed SAP system. You might need to modify these definitions if a SAP system for which you are receiving centralized alerts is also being monitored directly by another instance of the SAP agent. You do not want alerts that are reported under both systems. You can limit the centralized alert reports as follows:

- Use the /IBMMON/ITM\_CONFIG transaction to Maintain Managed Groups. Change the All CCMS alerts group. Remove the remote system from this list by editing the group definition to EXCLUDE the remote system identifier.
- Use the /IBMMON/ITM\_CONFIG transaction to Maintain Managed Groups. Delete the <remote\_SID> CCMS alerts group. For example, if the remote SAP system is QA2, this group name would be QA2 CCMS alerts.

Alternatively, you can use Centralized CCMS to report alerts from all SAP systems, but prevent alert reporting from each locally installed agent. Use the following steps to set up this configuration:

- Configure an instance of the SAP agent to monitor the Centralized CCMS system. Allow the agent to detect and report all alerts from all remote SAP systems.
- Configure an instance of the SAP agent to monitor each remote SAP system. Disable alert collection and reporting for these agent instances by using the /IBMMON/ITM\_CONFIG transaction to Select Monitor Sets and Monitors. Within this function, clear the check boxes for all monitors and save this configuration.

The SAP agent support for Centralized CCMS is used in a pure CCMS monitoring environment to view all alerts on a common console. Also, it can be used with its complete set of functions to provide situations, policies, and Take Action commands for the remote SAP systems.

### **Uninstalling the Advanced Business Application Programming (ABAP) transport from the SAP system**

If you choose to remove the SAP agent from your system, you must import Delete transport to the SAP system. Delete transport deletes the SAP agent dictionary objects and function modules.

#### **Before you begin**

If the SAP system is version 7.20 or later, before you import the delete transport, in your transport profile, you must add the following transport profile parameter: **tadirdeletions=true**. This transport profile parameter is available in tp version 375.57.68 and also in the R3trans version 6.14 release 700 or higher. For more information about removing transport requests from the SAP system, see [Deleting transport requests](#).

## Procedure

1. Go to the /ABAP directory on the product CD.
2. Copy the transport files into the SAP environment.
3. Copy the K711\_00xxx\_DELETE and R711\_00xxx\_DELETE files to the SAP Transport System data directory as follows:
  - a) Copy the K711\_00xxx\_DELETE file to the cofiles directory.
  - b) Copy the R711\_00xxx\_DELETE file to the data directory.
4. Run the following commands:
  - a) **tp addtobuffer ITMK711\_00xxx\_DELETE SID pf=\usr\sap\trans\bin\PROFILE\_NAME**
  - b) **tp import ITMK711\_00xxx\_DELETE SID client=nnn U16 pf=\usr\sap\trans\bin\PROFILE\_NAME**

Where:

### **SID**

Target SAP system ID

### **PROFILE\_NAME**

Name of the tp profile file

### **nnn**

Number for the target client where the agent is to run

## SAP instance customization

By default, all the instances of the SAP system are monitored and shown on the IBM Cloud App Management UI.

As an administrator, you choose which SAP instance you want to monitor. Also, as an administrator, you can turn off an SAP instance that you don't want to monitor.

The /IBMMON/ITM\_INSTANCE custom transaction links to the /IBMMON/ITM\_CONFIG transaction.

You select the **SAP Instances** option to view the available instances of the SAP server. Then, you select the instance that you want to monitor. These instances are displayed on the IBM Cloud App Management UI.

## Test Connection feature

The Test Connection feature verifies that you can connect your agent to the SAP system that is monitored.

Enter the parameters on the GUI to complete the test connection procedure. If you connect to the SAP system successfully, a success message is displayed. Alternatively, if the connection fails, a failure message is displayed.

## Enabling CCMS design

Computing Center Management System (CCMS) monitoring is enhanced to collect CCMS records that are in an open or closed state from the last sample period. You can configure the Sample period and by default it has a value of 3 minutes. However, you must ensure that the transport files that are referenced by the SAP agent and the Advanced Business Application Programming (ABAP) transport are the same version.

## Procedure

1. Log on to the SAP System.
2. Open the SE16 transaction and add the /IBMMON/ITM\_CNFG table name to the transaction.
3. Press **Enter** and then press **F8** to run the /IBMMON/ITM\_CNFG ABAP function module and to provide configurations for the ABAP program.
4. press **F5**. to create a new entry to which you add new configuration parameters.
5. In the **PARM NAME** field, enter ISNEWCCMSDESIGN and in the **VALUE CHAR** field, enter YES to create a new configuration parameter .



6. Click **Save**.

You can ignore the VALUE INT field.

### Modifying the threshold value of an alert

You can modify the **max ccms alert** threshold value that is associated with an alert. By default, the value is 1000, which means that you can view 1000 alerts in the IBM Cloud App Management. Older alerts are removed from the cache.

#### Procedure

Complete one of the following steps to modify the threshold value of an alert.

1. Follow the steps depending on your Operating system.
  - On Windows operating systems, open the <cancel home>\tmaitm6\KSAENV file.
  - On a Non-Windows operating systems, open the <cancel home>/config/sa.ini file.
2. Add the `MAX_CCMS_ALERT_THRESHOLD=< Value>` at the end of the file.

**Restriction:** The value must be greater than 100.

### Disabling CCMS design

You can disable Computing Center Management System (CCMS) design for SAP agent.

#### Procedure

1. Log on to the SAP System.
2. Open the SE16 transaction and add the /IBMMON/ITM\_CNFG table name to the transaction.
3. Press **Enter** and then press **F8** to run the /IBMMON/ITM\_CNFG ABAP function module and to provide configurations for the ABAP program.
4. Select and right-click **ISNEWCCMSDESIGN**, and then click **Delete**, to delete the existing entry.

## Configuring SAP HANA Database monitoring

You must configure the SAP HANA Database agent so that the agent can collect data of the SAP HANA database server that is being monitored.

### Before you begin

Review the hardware and software prerequisites, see [Software Product Compatibility Reports for SAP HANA Database agent](#).

Following are the prerequisites before you configure the SAP HANA Database agent:

1. Ensure to create users in all the databases (system and tenant) of the SAP HANA system with the following privileges:
  - Role: Monitoring
  - System privileges: Monitor Admin

The user name and password for the system and tenant databases must be the same.

2. When the switching between master to standby connectivity takes place on the SAP HANA Database agent system, the agent uses the hostname of Standby Server that needs to be resolved on the agent system. To resolve the hostname to an IP address, you need to add a mapping entry in host file of the machine on which the agent is installed.

**Note:** If you configure the agent by using Master Host, then enter the fully qualified host name or IP address of Master Host. If you configure the agent by using Stand by Host, then enter the fully qualified host name or IP address of Stand by Host. When you configure the agent through Stand by node, the Master node must be down along with the host machine.

### What to do next

Configure the SAP HANA Database agent on the operating system that you prefer.

- [“Configuring the agent on Windows systems” on page 486](#)
- [“Configuring the agent on Linux and AIX systems” on page 487](#)
- [“Configuring the agent by using the silent response file” on page 487](#)

## Configuring the agent on Windows systems

You can configure the SAP HANA Database agent on Windows systems.

### About this task

The SAP HANA Database agent is a multiple instance agent. You must create the first instance and start the agent manually.

### Procedure

To configure the agent on Windows systems, complete the following steps:

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click the **Monitoring Agent for SAP HANA Database** template, and then click **Configure agent**.
3. In the **Enter a unique instance name** field, type an agent instance name and click **OK**.

**Important:** The agent instance name must match the 3-digit HANA database system identifier (SID). For example, if the SID of the managed SAP HANA database is H01, enter H01 as the instance name.

4. In the **Monitoring Agent for SAP HANA Database** window, specify values for the following fields:

- **Instance Name**  
The default value for this field is identical to the value that you specified in the **Enter a unique instance name** field.
- **Server Name**  
The fully qualified host name or IP address of the SAP HANA server where the system database is installed.
- **Database Name**  
The name of the SAP HANA database.
- **Port Number**  
The SQL port number of the index server service on the system database of the SAP HANA database server.
- **HANA DB Administrator**  
The user name for accessing the SAP HANA database server.
- **HANA DB Administrator Password**  
The password for accessing the SAP HANA database server.
- **Confirm HANA DB Administrator Password**  
The password that is specified in the **HANA DB Administrator Password** field.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent on Linux and AIX systems

You can configure the SAP HANA Database agent on Linux and AIX systems so that the agent can collect data of the SAP HANA database server that is being monitored.

### Procedure

To configure the agent on Linux and AIX systems, complete the following steps:

1. On the command line, change the path to the agent installation directory. For example:

```
/opt/ibm/apm/agent/bin
```

2. Run the following command, where *instance\_name* is the name of the instance.

```
./sap_hana_database-agent.sh config instance_name
```

**Important:** The instance name must match the 3-digit HANA database system identifier (SID). If the SID of the managed SAP HANA database is H01, enter H01 as the instance name.

3. Enter 1 and press **Enter** when the command line displays the following message:  
Edit 'Monitoring Agent for SAP HANA Database' setting? [1=Yes, 2=No]
4. Specify values for the following agent parameters:
  - **Server Name**  
The fully qualified host name or IP address of the SAP HANA server where the system database is installed.
  - **Database Name**  
The name of the SAP HANA database.
  - **Port Number**  
The SQL port number of the index server service on the system database of the SAP HANA database server.
  - **HANA DB Administrator**  
The user name for accessing the SAP HANA database server.
  - **HANA DB Administrator Password**  
The password for accessing the SAP HANA database server.
  - **Confirm HANA DB Administrator Password**  
The password that is specified **HANA DB Administrator Password** field.
5. Run the following command to start the SAP HANA Database agent:

```
./sap_hana_database-agent.sh start instance_name
```

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent by using the silent response file

You can configure the SAP HANA Database agent using a silent response file so that the agent can collect data of the SAP HANA database server that is being monitored.

### Procedure

To configure the agent by using the silent response file, complete the following steps:

1. Edit the `sap_hana_silent_config.txt` file in editor of your preference and specify values for all the parameters.
  - For Windows systems, `sap_hana_silent_config.txt` file is available at `C:\IBM\APM\samples`.

- For Linux and AIX systems, `sap_hana_silent_config.txt` file is available at `/opt/ibm/apm/agent`.
2. On the command line, change the path to `install_dir`
  3. Run the following command:

- For Windows systems:

```
sap_hana_database-agent.bat config instance_name install_dir\samples
\sap_hana_silent_config.txt
```

- For Linux and AIX systems:

```
sap_hana_database-agent.sh config instance_name install_dir\samples
\sap_hana_silent_config.txt
```

4. Start the agent as follows:

- For Windows systems, go to the IBM Cloud Application Performance Management, Private window, right-click the agent instance that you created, and click **Start**.
- For Linux and AIX systems, run the following command:

```
./sap_hana_database-agent.sh start instance_name
```

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring SAP NetWeaver Java Stack monitoring

You must configure the SAP NetWeaver Java Stack agent so that the agent can collect resource monitoring data of the SAP NetWeaver Application Server that is being monitored.

### Before you begin

The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history”](#) on page 52.

Complete the following prerequisites before you configure the agent:

- Copy the following JAR files to the bin directory:
  - `sapj2eeclient.jar` (the SAP J2EE Engine client API that includes the JMX Adapter)
  - `logging.jar` (the logging library)
  - `com_sap_pj_jmx.jar` (the SAP-JMX library)
  - `exception.jar` (the SAP exception framework)

The bin directory is at the following path:

```
candle_home\TMAITM6_x64
candle_home\interp\sv\bin
```

**Important:** The JAR files are the same for all the supported operating systems. These files are available in the Diagnostics Agent patch or Software Update Manager (SUM).

### About this task

The SAP NetWeaver Java Stack agent is a multiple instance agent. You must create the first instance and start the agent manually.

- To configure the agent on Windows systems, you can use the GUI or the silent response file.

- To configure the agent on Linux or AIX systems, you can use the command line or the silent response file.

The directions that are mentioned in this topic are for the most current release of the agent, except as indicated. For information about how to check the version of an agent in your environment, see [Agent version](#).

## Configuring the agent on Windows systems

You can configure the agent on Windows operating systems by using the **IBM Performance Management** window.

### Before you begin

Ensure that the files, which are listed in the Before you begin section of the “[Configuring SAP NetWeaver Java Stack monitoring](#)” on page 488 topic, are available in the `bin` directory.

### About this task

The SAP NetWeaver Java Stack agent provides default values for some parameters. You can specify different values for these parameters.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Template** under the **Task/SubSystem** column, and click **Configure agent**.  
The **Monitoring Agent for SAP NetWeaver Java Stack** window opens.
3. In the **Enter a unique instance name** field, type an agent instance name and click **OK**.  
**Important:** The agent instance name must match the 3-digit SAP NetWeaver Java Stack system identifier (SID). For example, if the SID of the managed SAP NetWeaver Java Stack is P14, enter P14 as the instance name.
4. In the **Monitoring Agent for SAP NetWeaver Java Stack** window, specify values for the configuration parameters and click **OK**.  
For information about the configuration parameters, see “[Configuration parameters of the agent](#)” on page 491.
5. In the **IBM Performance Management** window, right-click the agent instance that you created and click **Start**.

### What to do next

- Log in to the Cloud App Management console to view the resource monitoring data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see “[Starting the Cloud App Management UI](#)” on page 176.

## Configuring the agent on Linux or AIX systems

To configure the agent on Linux or AIX systems, you must run the script and respond to prompts.

### Before you begin

Ensure that the files, which are listed in the "Before you begin" section of the “[Configuring SAP NetWeaver Java Stack monitoring](#)” on page 488 topic, are available in the `bin` directory.

### Procedure

1. On the command line, change the path to the agent installation directory.

```
Linux /opt/ibm/apm/agent/bin
```

**Linux** **UNIX** /opt/ibm/apm/agent/bin

2. Run the following command:

```
./sap_netweaver_java_stack-agent.sh config instance_name
```

where *instance\_name* is the name that you want to give to the instance.

**Important:** The agent instance name must match the 3-digit SAP NetWeaver Java Stack system identifier (SID). For example, if the SID of the managed SAP NetWeaver Java Stack is P14, enter P14 as the instance name.

3. When the command line displays the following message, type 1 and press Enter:

```
Edit 'Monitoring Agent for SAP NetWeaver Java Stack' setting? [1=Yes, 2=No]
```

4. When you are prompted, specify values for the configuration parameters.

For information about the configuration parameters, see [“Configuration parameters of the agent” on page 491](#)

5. Run the following command to start the agent:

```
./sap_netweaver_java_stack-agent.sh start instance_name
```

### What to do next

- Log in to the Cloud App Management console to view the resource monitoring data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring the agent by using the silent response file

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

### Before you begin

Ensure that the files, which are listed in the "Before you begin" section of the [“Configuring SAP NetWeaver Java Stack monitoring” on page 488](#) topic, are available in the bin directory.

### About this task

The silent response file contains the agent configuration parameters with default values defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters.

After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

1. In a text editor, open the `sap_netweaver_java_stack_silent_config.txt` file that is available at the following path, and specify values for the configuration parameters.

**Windows** C:\IBM\APM\samples

**Linux** **UNIX** /opt/ibm/apm/agent/samples

For information about the configuration parameters, see [“Configuration parameters of the agent” on page 491](#)

2. On the command line, change the path to `install_dir\bin`
3. Run the following command:

```
Windows sap_netweaver_java_stack-agent.bat config instance_name  
install_dir\samples\sap_netweaver_java_stack_silent_config.txt
```

```
Linux | UNIX ./sap_netweaver_java_stack-agent.sh config instance_name
install_dir\samples\sap_netweaver_java_stack_silent_config.txt
```

4. Start the agent.

```
Windows In the IBM Cloud Application Performance Management window, right-click the agent
instance that you created, and click Start. Alternatively, you can also run the following command:
sap_netweaver_java_stack-agent.bat start instance_name
```

```
Linux | UNIX Run the following command: ./sap_netweaver_java_stack-agent.sh
start instance_name
```

### What to do next

- Log in to the Cloud App Management console to view the resource monitoring data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuration parameters of the agent

When you configure the SAP NetWeaver Java Stack agent, you can change the default value of the parameters, such as **SAP\_NETWEAVER\_P4\_HOSTNAME** and **SAP\_NETWEAVER\_P4\_PORT**.

The following table contains detailed descriptions of configuration parameters of the SAP NetWeaver Java Stack agent. You must specify a value for all the fields because these fields are mandatory.

| Parameter name                           | Description                                                                                                                                               |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance Name</b>                     | The name of the instance. The default value for this field is identical to the value that you specified in the <b>Enter a unique instance name</b> field. |
| <b>SAP_NETWEAVER_P4_HOSTNAME</b>         | The host name or IP address of the SAP NetWeaver Application Server.                                                                                      |
| <b>SAP_NETWEAVER_P4_PORT</b>             | The P4 port number of the SAP NetWeaver Application Server.                                                                                               |
| <b>SAP_NETWEAVER_P4_USERNAME</b>         | The user name of the administrator for accessing the SAP NetWeaver Application Server.                                                                    |
| <b>SAP_NETWEAVER_P4_PASSWORD</b>         | The password of the administrator for accessing the SAP NetWeaver Application Server.                                                                     |
| Confirm <b>SAP_NETWEAVER_P4_PASSWORD</b> | The password that is specified for the <b>SAP_NETWEAVER_P4_PASSWORD</b> parameter.                                                                        |

## Configuring Skype for Business Server monitoring

When you install the Monitoring Agent for Skype for Business Server, the agent is in the unconfigured state. To start the agent, you need to configure it.

### Before you begin

- Review the hardware and software prerequisites. For the up-to-date system requirement information, see the [Software Product Compatibility Reports \(SPCR\) for the Skype for Business Server agent](#).
- Ensure that you are a domain user with administrator privileges and have access to all the remote servers that are listed in the Skype for Business Server topology. Use an existing domain user with administrator privileges, or create a new domain user and assign administrator privileges.

## About this task

You can configure the agent when the agent is in running or stopped state. The agent remains in the same state after configuration.

The product version and the agent version often differ. The directions here are for the most current release of this agent. For more information about how to check the version of an agent in your environment, see [Agent version command](#). For more information about the agent version list and what's new for each version, see the [“Change history” on page 52](#).

To configure the agent, you can either use the **IBM Performance Management** window or the silent response file.

## What to do next

After you configure the agent, you can change the user account from the local user to the domain user. For steps to change the user account, see [“Changing the user account” on page 494](#).

## Permissions and access rights for a non-administrator user

You can run the Monitoring Agent for Skype for Business Server as a non-administrator user but some functions are inaccessible in this case.

### Registry Permissions

To create a non-administrator user, create a new user and set up registry permissions for the new user as follows:

- Full access to the KEY\_LOCAL\_MACHINE\SOFTWARE\IBMMonitoring
- Full access to the CANDLE\_HOME directory

The non-administrator user must be a member of the Performance Monitor Users and Performance Log Users group. If you define these permissions for a non-administrator user, data is displayed for all the Perfmon-based attribute groups.

### Viewing attribute groups' data collected from Database

If you want to view data of an attribute group, which is collected from database, you must set up the following permissions for the non-administrator user.

- The non-administrator user account that is used to run the Skype for Business Server agent must have the Debug Program permission to add a debugger to any process.

By default, the Debug Program permission is assigned only to the administrator and Local System accounts. To grant the Debug Program permission, you must complete the following steps on the Skype for Business Server:

1. Click **Start > Administrative Tools > Local Security Policy**. The **Local Security Settings** window opens.
  2. Expand **Local Policies** and click **User Rights Assignment**. The list of user rights opens.
  3. Double-click **Debug Programs policy**. The **Debug programs Properties** window opens.
  4. Click **Add User or Group**. The **Select Users or Groups** window opens.
  5. In **Enter the object names to select** field, enter the user account name to whom you want to assign permissions and click **OK**.
  6. Click **OK**.
- Grant Log on as Service permission

To grant the Log-on as service permission, you must complete the following steps on the Skype for Business Server:

1. Click **Start > Administrative Tools > Local Security Policy**. The **Local Security Settings** window opens.



2. Expand **Local Policies** and click **User Rights Assignment**. The list of user rights opens.
3. Double-click **Log-on** as service policy. The **Log-on as service Properties** window opens.
4. Click **Add User or Group**. The **Select Users or Groups** window opens.
5. In **Enter the object names to select** field, enter the user account name to whom you want to assign permissions and click **OK**.
6. Click **OK**.

The Availability attribute group shows the data for users who are members of the Administrators group.

## Configuring the agent on Windows systems

You can configure the Skype for Business Server agent on Windows operating systems by using the **IBM Performance Management** window. After you update the configuration values, you must start the agent to save the updated values.

### About this task

Configure the agent when the agent is running or stopped state. The agent remains in the same state after configuration.

The Skype for Business Server agent provides default values for some parameters. You can specify different values for these parameters.

### Procedure

1. Click **All Programs > Start > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Skype for Business Server agent** and click **Configure agent**.
3. In the Skype for Business Server agent window, complete the following steps:
  - a) On the **SQL Configuration for Skype for Business Topology** tab, to connect to the Microsoft Lync Server or Skype for Business Server Central Management Store, specify values for the configuration parameters, and click **Next**.
 

**Note:** You can skip this tab, as SQL Configuration for Skype for Business Topology is not applicable for IBM Cloud Application Management.

**Important:** Synthetic transaction configuration is optional. If you require the synthetic transaction data, enter the configuration parameters on the **Setup Information** and **Scheduler Configuration** tabs.
  - b) On the **Administrator Login Credentials** tab, enter the administrator credentials, and click **Next**.
  - c) On the **Setup Information** tab, to run commands for the synthetic transactions, enter the values for the configuration parameters and click **Next**.
  - d) On the **Scheduler Configuration** tab, to schedule the synthetic transactions, enter the values for the configuration parameters and click **Next**.
  - e) On the **SQL Server Configuration for Skype for Business Monitoring Role** tab, to connect to the Microsoft Lync Server or Skype for Business Server monitoring role, enter the values for the configuration parameters and click **Next**.

For more information, see [“Configuration parameters for the agent”](#) on page 495.

4. In the **IBM Performance Management** window, right-click **Monitoring Agent for Skype for Business Server** and click **Start**.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring the agent by using the silent response file

You can configure the Skype for Business Server agent by using silent response file. It contains the agent configuration parameters with default values that are defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters.

### Before you begin

If you want to modify the default configuration parameters, edit the response file.

### About this task

You can configure the agent when the agent is running or stopped. The agent remains in the same state after configuration. For example, if the agent is running, it remains in the running state after configuration.

### Procedure

To configure the Skype for Business Server agent, complete the following steps:

1. On the command prompt, change the path to the directory that contains the `skype_for_business_server-agent.bat` file.
2. Run the following command:

```
skype_for_business_server-agent.bat config absolute path to the response file
```

For information about the configuration parameters, see [“Configuration parameters for the agent” on page 495](#).

3. Optional: If the agent is in stopped state, start the agent.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Changing the user account

After you configure the Skype for Business Server agent, you can change the user account from the local user to the domain user.

### About this task

By default, the Skype for Business Server agent runs under the local user account. To collect data from the remote servers, the agent must run under the domain user.

### Procedure

1. Run the following command to verify which user ID is being used for starting the agent:

```
install_dir\InstallITM\KinCinfo.exe -r
```

2. If the monitoring agent is started with a user ID that does not belong to the Administrator group, stop the agent.
3. Open the **Manage Monitoring Services** window.
4. Right-click on agent instance and click **Change Startup**.
5. Enter the fully qualified user ID as `<Domain\User ID>` and password.
6. Start the Skype for Business Server agent.

## Configuration parameters for the agent

When you configure the Skype for Business Server agent, you can change the default values of the configuration parameters, such as the database server name, database instance name, database name, and others.

The following table contains descriptions of the configuration parameters for the Skype for Business Server agent.

**Note:** Out of all the fields, the Pool Fully Qualified Domain Name field is mandatory in following table.

| <i>Table 66. Names and descriptions of the configuration parameters for the agent</i> |                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameter name</b>                                                                 | <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                            |
| Database Server Name                                                                  | <ul style="list-style-type: none"> <li>• <b>SQL Configuration for Skype for Business Topology</b> tab: The name of the database server where the Lync or Skype for Business Server Central Management Store is installed.</li> <li>• <b>SQL Server Configuration for Skype for Business Monitoring Role</b> tab: The name of the database server where the monitoring role is installed.</li> </ul> <p>Example is PS6877.</p> |
| Database Instance Name                                                                | <ul style="list-style-type: none"> <li>• <b>SQL Configuration for Skype for Business Topology</b> tab: The default instance.</li> <li>• <b>SQL Server Configuration for Skype for Business Monitoring Role</b> tab: The name of the database instance where the monitoring role is installed.</li> </ul>                                                                                                                      |
| Database Name                                                                         | The name of the database.                                                                                                                                                                                                                                                                                                                                                                                                     |
| Database User ID                                                                      | The user ID of the database. This user must have access to the required Microsoft SQL Server instance. This user can or cannot be an Active Directory user.                                                                                                                                                                                                                                                                   |
| Database Password                                                                     | The password of the database where the monitoring role is installed.                                                                                                                                                                                                                                                                                                                                                          |
| Username (Example: skype\administrator)                                               | The user ID of the administrator. This user must be a domain user with administrator privileges and access to all the remote servers that are listed in the Lync or Skype for Business Server topology. The credentials of this user are also used in Synthetic Transaction feature. So, this user must be authorized to create Windows Schedule in Task Scheduler and run Synthetic Transaction Commands.                    |
| Password                                                                              | The login password of administrator.                                                                                                                                                                                                                                                                                                                                                                                          |
| Confirm Domain Password                                                               | Enter the same password that you specified in the Domain Password field.                                                                                                                                                                                                                                                                                                                                                      |
| Pool FQDN                                                                             | The fully qualified domain name (FQDN) of Skype Pool for which you run the synthetic commands.                                                                                                                                                                                                                                                                                                                                |
| Geographic Location                                                                   | The geographic location of the production system.                                                                                                                                                                                                                                                                                                                                                                             |
| Test Users1 (for example, user1@skype.com)                                            | <p>First Username that can be used to run Synthetic Transaction cmdlets. Format for username is SAMAccountName@domain.com</p> <p><b>Restriction:</b> Do not provide Sip Address.</p>                                                                                                                                                                                                                                          |
| Test User1 PWD                                                                        | The password of Test User1.                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 66. Names and descriptions of the configuration parameters for the agent (continued)

| Parameter name                            | Description                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Confirm Test User1 PWD                    | Enter the same password that you specified in the <b>Test User1 PWD</b> field.                                                                                                                                                                                                                                                                                                             |
| Test User2 (for example, user2@skype.com) | Second Username that can be used to run Synthetic Transaction cmdlets. Format for username is SAMAccountName@domain.com<br><br><b>Restriction:</b> Do not provide Sip Address.                                                                                                                                                                                                             |
| Test User2 PWD                            | The password of Test User2.                                                                                                                                                                                                                                                                                                                                                                |
| Confirm Test User2 PWD                    | Enter the same password that you specified in the <b>Test User2 PWD</b> field.                                                                                                                                                                                                                                                                                                             |
| Use Agent Configuration Values            | Keep this field enabled if you want to run synthetic commands by using all fields provided in configuration window.<br>Disable to use values set by New-CsHealthMonitoringConfiguration. If disabled, the value of <b>Pool1 FQDN</b> is used as identity for Get-CsHealthMonitoringConfiguration. Make sure to provide valid test user credentials to run <b>Test -CsMcxP2PIM</b> command. |
| Frequency                                 | The frequency of the scheduled utility that fetches the data of synthetic transactions. The frequency has the following values: <ul style="list-style-type: none"> <li>• Daily (DAY_FREQUENCY)</li> <li>• Weekly (WEEK_FREQUENCY)</li> <li>• Monthly (MONTHLY_FREQUENCY)</li> </ul>                                                                                                        |
| Collection Hour                           | The hour part of the time-stamp, in the 24-hour clock format that you select to schedule the utility.                                                                                                                                                                                                                                                                                      |
| Collection Minute                         | The minutes part of the time-stamp that you select to schedule the utility.                                                                                                                                                                                                                                                                                                                |
| Start Date (YYYY-MM-DD)                   | The time when the scheduler is activated.                                                                                                                                                                                                                                                                                                                                                  |
| End Date (YYYY-MM-DD)                     | The time when the scheduler is deactivated.                                                                                                                                                                                                                                                                                                                                                |

## 2019.4.0.2 Configuring Sterling Connect Direct monitoring

You must configure the Sterling Connect Direct agent so that the agent can collect data from the Connect Direct nodes to monitor the availability and performance of Connect Direct nodes. You can configure the agent on Windows and Linux systems.

### Before you begin

Ensure that the system requirements for the Sterling Connect Direct agent are met in your environment.

### About this task

The Sterling Connect Direct agent is a multiple instance agent. You must create the first instance and start the agent manually.

- To configure the agent on Windows systems, you can use the IBM Cloud Application Performance Management window or the silent response file.

- To configure the agent on Linux systems, you can run the script and respond to prompts, or use the silent response file.

### 2019.4.0.2 Configuring the agent on Windows systems

To configure the agent on Windows operating systems, you can use the IBM Cloud Application Performance Management window. After you update the configuration values, start the agent to apply the updated values.

#### About this task

The Sterling Connect Direct agent provides default values for some parameters. You can specify different values for these parameters.

#### Procedure

To configure the agent on Windows operating systems, follow these steps:

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Template** in the **Task/SubSystem** column, and click **Configure agent**.
3. In the **Enter a unique instance name** field, type an agent instance name and click **OK**.
 

**Note:** Limit the length of agent instance name. Preferably in the range of 7 - 10 characters.
4. In the **Monitoring Agent for Sterling Connect Direct** window, go to the **Connect Direct Server Details** tab, specify values for the configuration parameters and click **OK**. For more information about the configuration parameters, see [“Configuration parameters of the agent”](#) on page 499.
5. Click **Next**.
6. On the Java Parameters tab, keep default values and click **Next**.
7. On the Java API Client Configuration tab, click **OK**.
8. In the **IBM Performance Management** window, right-click the agent instances that you created and click **Start** to start the agent.

#### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

### 2019.4.0.2 Configuring the agent on Linux systems

To configure the agent on Linux operating systems, you must run the script and respond to prompts.

#### Procedure

To configure the agent on Linux operating systems, follow these steps:

1. On command line, change the path to the agent installation directory.  
For example, `/opt/ibm/apm/agent/bin`
2. Run the following command:  
**`/sterling_connect_direct-agent.sh config instance_name`**  
Where *instance\_name* is the name you want to give to the agent instance.
3. Command line displays the following message:  
Edit 'Monitoring Agent for Sterling Connect Direct' setting? [1=Yes, 2=No]
4. Enter 1 to edit the settings.
5. When you are prompted to enter a value for the following parameters, press Enter to accept the default value, or specify a different value and press enter.

- Instance name
- Server name
- Server port
- User name
- Password
- JAVA home
- Java trace level (Default value is `ERROR`.)
- JVM arguments
- Class path for external JAR

For more information about the configuration parameters, see [“Configuration parameters of the agent” on page 499](#).

6. Run the following command to start the agent:

```
./sterling_connect_direct-agent.sh start instance_name
```

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## 2019.4.0.2 Configuring the agent by using the silent response file

Use the silent response file to configure the agent without responding to prompts when you run the configuration script. You can use the silent response file for configuring the agent on both Windows and Linux systems.

### About this task

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the silent configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode. After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

To configure the agent by using the silent response file, follow these steps:

1. In a text editor, open the silent response file that is available at the following path:  
`install_dir/samples/sterling_connect_direct_silent_config.txt`

Where `install_dir` is the installation directory of the Sterling Connect Direct agent.

2. In the response file, specify a value for the following parameters:

- For the **Server Name** parameter, specify the hostname or IP of Sterling Connect Direct server that you want to monitor. Otherwise, retain the default value as `localhost`.
- For the **User name** parameter, enter the user name.
- For the **Password** parameter, enter the password.

3. Save and close the response file, and run the following command to update the agent configuration settings:

```
Linux | UNIX ./sterling_connect_direct-agent.sh config  
<Instance_name>install_dir
```

```
Windows ./sterling_connect_direct-agent.bat config <Instance_name>install_dir
```

Where *instance\_name* is the name that you want to give to the instance, and *install\_dir* is the installation directory of Sterling Connect Direct agent.

**Important:** Be sure to include the absolute path to the silent response file. Otherwise, no agent data is displayed in the dashboards.

4. Start agent by using the following command:

**Linux** | **UNIX** Run the `install_dir\bin\sterling_connect_direct-agent.sh start` command.

**Windows** Right-click **Monitoring Agent for Sterling Connect Direct** and then click **Start**.

### What to do next

Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For more information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## 2019.4.0.2 Configuration parameters of the agent

When you configure the Sterling Connect Direct agent, you can set the values for configuration parameters.

The following table contains detailed description of the configuration parameters for the Sterling Connect Direct agent.

*Table 67. Name and description of the configuration parameters*

| Parameter name              | Description                                                                                                                   | Mandatory field |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Instance Name               | The default value for this field is identical to the value that you specify in the <b>Enter a unique instance name</b> field. | Yes             |
| Server name                 | The hostname or IP of Sterling Connect Direct server.                                                                         | Yes             |
| Server port                 | The Port of Sterling Connect Direct server. Default value for Sterling Connect Direct is 1363.                                | Yes             |
| Username                    | The username to connect to Sterling Connect Direct server.                                                                    | Yes             |
| Password                    | The password to connect to Sterling Connect Direct server.                                                                    | Yes             |
| Java Home                   | Path to the folder where Java is installed.                                                                                   | No              |
| Java trace level            | The trace level used by Java providers. Default value for Sterling Connect Direct is Error.                                   | Yes             |
| JVM arguments               | An optional list of arguments to the Java virtual machine.                                                                    | No              |
| Class path for external JAR | The path for JAR required by Java API data provider that is not included with the agent.                                      | No              |

## 2019.4.0.2 Configuring Sterling File Gateway monitoring

The Monitoring Agent for Sterling File Gateway monitors the IBM® Sterling File Gateway application by using the business-to-business (B2B) REST APIs and file gateway database. You must configure the Sterling File Gateway agent so that the agent can collect data from the data sources and monitor the statistics and health of the Sterling File Gateway application. You can configure the agent on Windows and Linux systems.

### Before you begin

- Ensure that the system requirements for the Sterling File Gateway agent are met in your environment.

- Ensure that the B2B REST APIs are installed on your file gateway node. For more information about the B2B REST API installation, see [“Installing the B2B REST API” on page 500](#).

### About this task

The Sterling File Gateway agent is a multiple instance agent. You must create the first instance and start the agent manually.

- To configure the agent on Windows systems, you can use the **IBM Performance Management** window or the silent response file.
- To configure the agent on Linux systems, you can run the script and respond to prompts, or use the silent response file.

## 2019.4.0.2 Installing the B2B REST API

You can install and configure the business-to-business (B2B) REST APIs on your Sterling File Gateway node. The B2B REST APIs are available in the B2B Integrator installer (V5.2.6.2).

### Procedure

To install the B2B REST API, follow these steps:

1. Go to the `<install_dir>/bin` directory.

Where, *install\_dir* is the agent installer directory for the B2B integrator.

2. Run the following command:

- **Linux** `./InstallService.sh/install_dir/bin/b2bAPIs_10000602.jar`

Where, *<install\_dir>* is the location where you extracted the media file content.

- **Windows** `./InstallService.cmd/install_dir/bin/b2bAPIs_10000602.jar`

Where, *<install\_dir>* is the B2B installer folder.

## 2019.4.0.2 Configuring the Sterling File Gateway agent on Windows systems

You can configure the Sterling File Gateway agent on Windows operating systems by using the **IBM Cloud Application Performance Management** window. After you update the configuration values, you must start the agent to apply the updated values.

### About this task

The Sterling File Gateway agent provides default values for some parameters. You can specify different values for these parameters.

### Procedure

To configure the agent on Windows operating systems, follow these steps:

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Monitoring Agent for Sterling File Gateway**, and then click **Configure agent**.

**Remember:** After you configure the agent for the first time, the **Configure agent** option is not available. To configure the agent again, click **Reconfigure**.

3. In the Sterling File Gateway agent window, complete the following steps:
  - a) Enter a unique name for the Sterling File Gateway agent instance, and click **OK**.
  - b) On the **B2B API Details** tab, specify values for the configuration parameters, and then click **Next**.
  - c) On the **Database Details** tab, specify values for the configuration parameters, and then click **Next**.
  - d) On the **Java API** tab, specify values for the configuration parameters, and then click **OK**.



For more information about the configuration parameters in each tab of the Sterling File Gateway agent window, see the following topics:

- “[Configuration parameters for the B2B API details](#)” on page 504
  - “[Configuration parameters for database details](#)” on page 505
  - “[Configuration parameters for the Java API](#)” on page 505
4. In the **IBM Performance Management** window, right-click **Sterling File Gateway agent**, and then click **Start**.

## 2019.4.0.2 Configuring the Sterling File Gateway agent on Linux systems

You can run the configuration script and respond to prompts to configure the Sterling File Gateway agent on the Linux operating systems.

### Procedure

To configure the agent on Linux operating systems, follow these steps:

1. Go to command line and run the following command:

```
<install_dir>/bin/sterling_file_gateway-agent.sh config instance_name
```

Where, *instance\_name* is the name that you want to give to the instance and *install\_dir* is the agent installation directory path.

2. You are prompted to provide values for all the mandatory configuration parameters. You can modify the default values of configuration parameters.

For more information about the configuration parameters, see the following topics:

- “[Configuration parameters for the B2B API details](#)” on page 504
  - “[Configuration parameters for database details](#)” on page 505
  - “[Configuration parameters for the Java API](#)” on page 505
3. To start the agent, run the following command:

```
<install_dir>/bin/sterling_file_gateway-agent.sh start instance_name
```

## 2019.4.0.2 Configuring Sterling File Gateway agent by using the silent response file

Use the silent response file to configure the agent without responding to prompts when you run the configuration script. You can use the silent response file for configuring the agent on both Windows and Linux systems.

### About this task

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the silent configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode. After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

### Procedure

- To configure the agent by using the silent response file, follow these steps:
  - a) In a text editor, open the `sterling_file_gateway_silent_config.txt` file that is available at the following path:

```
- Linux install_dir/samples/sterling_file_gateway_silent_config.txt
```

For example, `/opt/ibm/apm/agent/samples/sterling_file_gateway_silent_config.txt`

- **Windows** `install_dir\samples\sterling_file_gateway_silent_config.txt`

For example, **C:\IBM\APM\samples\sterling\_file\_gateway\_silent\_config.txt**

- b) In the `sterling_file_gateway_silent_config.txt` file, specify values for all the mandatory parameters. You can also modify the default values of other parameters.

For more information about the configuration parameters, see the following topics:

- [“Configuration parameters for the B2B API details” on page 504](#)
- [“Configuration parameters for database details” on page 505](#)
- [“Configuration parameters for the Java API” on page 505](#)

- c) Save and close the `sterling_file_gateway_silent_config.txt` file, and run the following command:

- **Linux** `install_dir/bin/sterling_file_gateway-agent.sh config install_dirinstall_dir/samples/sterling_file_gateway_silent_config.txt`

For example, **/opt/ibm/apm/agent/bin/sterling\_file\_gateway-agent.sh config instance\_name /opt/ibm/apm/agent/samples/sterling\_file\_gateway\_silent\_config.txt**

- **Windows** `install_dir/bin/sterling_file_gateway-agent.bat config instance_name install_dir/samples/sterling_file_gateway_silent_config.txt`

For example, **C:\IBM\APM\bin\sterling\_file\_gateway-agent.bat config instance\_name C:\IBM\APM\samples\sterling\_file\_gateway\_silent\_config.txt**

Where, *instance\_name* is the name that you want to give to the instance and *install\_dir* is the path where the agent is installed.

**Important:** Ensure that you include the absolute path to the silent response file. Otherwise, the agent data is not shown in the dashboards.

- d) Run the following command to start the agent:

- **Linux** `install_dir/bin/sterling_file_gateway-agent.sh start instance_name`

For example, **/opt/ibm/apm/agent/bin/sterling\_file\_gateway-agent.sh start instance\_name**

- **Windows** `install_dir\bin\sterling_file_gateway-agent.bat start instance_name`

For example, **C:\IBM\APM\bin\sterling\_file\_gateway-agent.bat start instance\_name**

## 2019.4.0.2 Configuring agent environment variables for the data provider on Linux

You can configure the Sterling File Gateway agent environment variables for the data provider on Linux operating systems.

### About this task

The Sterling File Gateway agent provides environment variables that you can configure for the data provider.

### Procedure

To configure agent environment variables for the data provider on Linux, follow these steps:

1. Go to the `<install_dir>/agent/config` directory.

2. Open the .fg.environment file in an editor and edit the environment variables.

For more information about the agent environment variables that you can configure, see [“Environment variables for the data provider”](#) on page 503.

### 2019.4.0.2 **Configuring agent environment variables for the data provider on Windows**

You can configure the Sterling File Gateway agent environment variables for the data provider on Windows operating systems by using the **IBM Performance Management** window.

#### **About this task**

The Sterling File Gateway agent provides environment variables that you can configure for the data provider.

#### **Procedure**

To configure agent environment variables for the data provider on Windows, follow these steps:

1. Click **Start > All Programs > > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click the agent instance and click **Advanced > Edit ENV File** and edit default values for the environment variables.

For more information about the agent environment variables that you can configure, see [“Environment variables for the data provider”](#) on page 503.

### 2019.4.0.2 **Environment variables for the data provider**

After you configure the Sterling File Gateway agent, you can modify some threshold duration values related to the agent data collection. You can specify these values in the agent environment file.

The following table contains detailed description of the environment variables for the data provider.

| <b>Parameter name</b>                                                                                                                   | <b>Description</b>                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Collection duration for files transfer (in hours)<br><b>(KFG_FILE_ARRIVED_INTERVAL)</b>                                                 | The duration in hours for which the agent collects data for file transfers. The default value is 24 hours.                                                                                                                                                                                                       |
| Collection intervals for files transfer activities that are displayed as a line chart (in hours)<br><b>(KFG_FILE_ACTIVITY_INTERVAL)</b> | The duration in hours for which the agent collects data for file transfer activities. The default value is 1 hour.<br><br>For example, the agent collects file transfer activities that occurred in last 1 hour. This data is visible in terms of line charts on the instance page. The default value is 1 hour. |
| Threshold interval for inactive partners (in days)<br><b>(KFG_INACTIVE_PARTNERS_INTERVAL)</b>                                           | The threshold duration when the partner is inactive or not received or uploaded any file. The default value is 10 days.<br><br>For example, if any partner that does not receive or transfer a file in last 10 days, displays as "Inactive" on the agent.                                                        |
| Maximum number of data provider log files<br><b>(KFG_LOG_FILE_MAX_COUNT)</b>                                                            | The maximum number of log files that the data provider creates before it overwrites the previous log files. The default value is 10.                                                                                                                                                                             |
| Maximum size in KB of each data provider log<br><b>(KFG_LOG_FILE_MAX_SIZE)</b>                                                          | The maximum size in KB that a data provider log file must reach before the data provider creates a new log file. The default value is 5190 KB.                                                                                                                                                                   |

Table 68. Name and description of the environment variables for the data provider (continued)

| Parameter name                                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Level of detail in data provider log (KFG_LOG_LEVEL)      | <p>The level of details that are included in the log file that the data provider creates. The default value is 4 (information). The following values are valid:</p> <ul style="list-style-type: none"> <li>• 1 (Off): No messages are logged.</li> <li>• 2 (Severe): Only errors are logged.</li> <li>• 3 (Warning): All errors and messages that are logged at the severe level and potential errors that might result in undesirable behavior.</li> <li>• 4 (Info): All errors and messages that are logged at the warning level and high-level informational messages that describe the state of the data provider when it is processed.</li> <li>• 5 (Fine): All errors and messages that are logged at the information level and low-level informative messages that describe the state of the data provider when it is processed.</li> <li>• 6 (Finer): All errors and messages that are logged at the fine level plus detailed informative messages, such as performance profiling information and debug data. Selecting this option can adversely affect the performance of the monitoring agent. This setting is intended only as a tool for problem determination along with the IBM support staff.</li> <li>• 7 (Finest): All errors and messages that are logged at the Fine level and the most detailed informative messages that include low-level programming messages and data. Selecting this option might adversely affect the performance of the monitoring agent. This setting is intended only as a tool for problem determination along with the IBM support staff.</li> <li>• 8 (All): All errors and messages are logged.</li> </ul> |
| Fetching events for all file transfers (KFG_ALL_FGEVENTS) | <p>The flag for fetching events for all file transfers. The valid values are, Yes or No. The default value is No. If the value is set to No, then agent fetches events for failed file transfers for a user configurable duration. If the value is set to Yes, then the agent fetches events for all file transfers for a user configurable duration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

### 2019.4.0.2 Configuration parameters for the B2B API details

When you configure the Sterling File Gateway agent, you must specify values of the configuration parameters for the business-to-business (B2B) API details.

The following table contains detailed description of the configuration parameters for the B2B API details.

| Table 69. Name and description of the configuration parameters for the B2B API details |                                                                                                                                                                                                                                                                               |
|----------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Parameter name                                                                         | Description                                                                                                                                                                                                                                                                   |
| Instance Name (KFG_Instance_Name)                                                      | <p>The name of the instance.</p> <p><b>Restriction:</b> The Instance Name field displays the name of the instance that you specify when you configure the agent for the first time. When you configure the agent again, you cannot change the instance name of the agent.</p> |
| Server Name (KFG_API_SERVICES_Node_ADDRESS)                                            | <p>The hostname or IP address of the B2B API service.</p>                                                                                                                                                                                                                     |

Table 69. Name and description of the configuration parameters for the B2B API details (continued)

| Parameter name                           | Description                                                                    |
|------------------------------------------|--------------------------------------------------------------------------------|
| Server Port<br>(KFG_API_SERVICES_PORT)   | The port of the B2B API.                                                       |
| User Name<br>(KFG_API_SERVICES_USERNAME) | A user name to connect to the B2B API service.                                 |
| Password<br>(KFG_API_SERVICES_PASSWORD)  | The password for the user name that you use to connect to the B2B API service. |

### 2019.4.0.2 Configuration parameters for database details

When you configure the Sterling File Gateway agent, you must specify values of the configuration parameters for the database details.

The following table contains detailed description of the configuration parameters for the database details.

Table 70. Name and description of the configuration parameters for the database details

| Parameter name                                | Description                                                               |
|-----------------------------------------------|---------------------------------------------------------------------------|
| Database Server Name<br>(KFG_DB_Node_ADDRESS) | The host name or IP address of the Sterling File Gateway database server. |
| Database User (KFG_DB_USERNAME)               | The name of the database user.                                            |
| Database Password (KFG_DB_PASSWORD)           | The password of the database.                                             |
| Database Port (KFG_DB_PORT)                   | The port of the database.                                                 |
| Database Type (KFG_DB_TYPE)                   | The type of the database.                                                 |

### 2019.4.0.2 Configuration parameters for the Java API

When you configure the Sterling File Gateway agent, you must specify values of the configuration parameters for the Java API.

The following table contains detailed description of the configuration parameters for the Java API.

Table 71. Name and description of the configuration parameters for the Java API

| Parameter name                                     | Description                                                                           |
|----------------------------------------------------|---------------------------------------------------------------------------------------|
| Class path for the external JAR<br>(KFG_CLASSPATH) | The database driver JAR path that you want to specify for the corresponding database. |

## Configuring Sybase Server monitoring

The Sybase agent offers a central point of management for distributed databases. It collects the required information for database and system administrators to examine the performance of the Sybase Server system, detect problems early and prevent them. Database and system administrators can set the required threshold levels and flags to trigger alerts when the system reaches these thresholds. You must configure the Monitoring Agent for Sybase Server to monitor Sybase Server.

### Before you begin

- The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history”](#) on page 52.

## About this task

The Sybase agent is a multiple instance agent, you must configure and start each agent instance manually.

## Procedure

1. Configure the monitoring agent.
  - [“Configuring the agent by using the command line interface” on page 507](#)
  - [“Configuring the agent by using the silent response file” on page 509](#)
2. Start and stop the monitoring agent by using agent command **sybase-agent**.  
For more information about the **sybase-agent**, see [“Using agent commands” on page 226](#).
3. Connect the monitoring agent to the Performance Management server by using the command **agent2server**.  
For more information about the **agent2server**, see [“Using agent commands” on page 226](#).

## Granting permissions

You must grant permissions to the user ID that is used to monitor the Sybase Server.

### Before you begin

Install the Sybase agent.

You must have the database administrator role to grant permissions.

### About this task

The user ID that is used by the monitoring agent must have access to Sybase tables and the installed monitor tables.

You can perform the following tasks:

- Create a user ID for the monitoring agent.
- Grant permission to the new user ID and the installed monitor tables.

If you are not running the Sybase agent as root user, make sure that the user ID belongs to the Sybase group and has read-access to the Sybase log files.

## Procedure

1. Enter the command for the operating system you are using.

- **Windows**

```
cd install_dir\tmaitm6\SQLLIB
```

- **UNIX**

```
cd install_dir/misc
```

Where, *install\_dir* is the home directory where the Sybase agent is installed.

2. Use the **isql** command to log in to the Sybase Server as user sa.
3. Run the following command to configure the ID that is used by the Sybase agent to communicate with Sybase Server:

```
1>sp_addlogin user_name, password 2>g
```

Where:

- *user\_name* is the user ID. By default, it is `tivoli`.

If the user ID is not `tivoli`, edit the `koygrant.sql` file and change the `tivoli` to the correct user ID.

- *password* is the password of the user.

**Note:**

Location of the `koygrant.sql` file:

- **Windows** `install_dir\tmaitm6\SQLLIB`  
Where *install\_dir* is the home directory where the Sybase agent is installed.
- **UNIX** `/opt/ibm/apm/agent/misc/`

4. Run the following command to grant permission to the tables in the database:

```
isql -U sa -P password -S servername -i koygrant_filepathkoygrant.sql
```

Where:

- *password* is the password of user `sa`.
- *servername* is the database server name.
- *koygrant\_filepath* is at the following location:

**Note:**

- **Windows** `\opt\ibm\apm\agent\misc\`
- **UNIX** `/opt/ibm/apm/agent/misc/`

5. Run the following command to create proxy tables that are used for the installed monitor tables:

```
isql -U sa -P password -S servername  
-i $SYBASE/ASE-12_5/scripts/installmontables
```

Where:

- *password* is the password of user `sa`.
- *servername* is the database server name.

## What to do next

When the permissions are successfully granted, you can configure the monitoring agent.

## Configuring the agent by using the command line interface

You can configure the Monitoring Agent for Sybase Server by using the command line interface.

### Before you begin

The Sybase agent does not support remote configuration. Hence, you need to ensure that the Sybase Server is installed on the same host where the Sybase agent is installed.

The Sybase agent supports Sybase Server version 15.7 and 16.0 only.

The user ID that is used to connect to the database server is created.

### About this task

The Sybase agent is a multiple instance agent, you must configure and start each agent instance manually.

### Procedure

1. Run the following command to configure the agent.

- **Windows**

```
install_dir\bin\sybase-agent.bat config instance_name
```

- **UNIX**

```
install_dir/bin/sybase-agent.sh config instance_name
```

Where:

- *install\_dir* is the agent installation directory.
  - *instance\_name* is the Sybase Server instance name.
2. When you are prompted to provide values for the following parameters, press Enter to accept the default value, or specify a value and press Enter.
- a) For the Home Directory parameter, enter the path of the Sybase Server home directory path.
    - **Windows**  
The example of Home Directory is \opt\sybase.
    - **UNIX**  
The example of Home Directory is /opt/sybase.
  - b) For the ASE Directory parameter, enter the path of the database server ASE.
    - **Windows**  
The example of ASE Directory is \opt\sybase\ASE-12\_5.
    - **UNIX**  
The example of ASE Directory is /opt/sybase/ASE-12\_5.
  - c) For the Open Client Directory parameter, enter the Sybase open client installation location.
    - **Windows**  
The example of Open Client Directory is \opt\sap\ocs-16\_0.
    - **UNIX**  
The example of Open Client Directory is /opt/sap/ocs-16\_0.
  - d) For the USER ID parameter, enter the user ID that is used by the monitoring agent to connect to the Sybase Server.  
The default USER ID is `tivoli`.
  - e) For the PASSWORD parameter, enter the password of the user ID that is used by the monitoring agent to connect to the Sybase Server.
  - f) For the VERSION parameter, enter the Sybase Server version.  
The Sybase agent supports Sybase Server versions 15.7 and 16.0 only.
  - g) For the ERROR LOG FILE parameter, enter the fully qualified file name of the error log file for the Sybase Server.
    - **Windows**  
The example of the ERROR LOG FILE is \opt\sap\ASE-16\_0\install\servername.log.
    - **UNIX**  
The example of the ERROR LOG FILE is /opt/sap/ASE-16\_0/install/servername.log.Where *servername* is the Sybase Server name.
  - h) For the EXTENDED parameter, enter the extended parameter that is used by support to exclude certain cursor execution. Optionally, press Enter without specifying any values to execute all cursors.  
The options for the EXTENDED parameter are `DBD2`, `DBD15`, `KOYSEGD`.



- DBD2 will exclude cursor execution for datasets Sybase\_Database\_Detail and Sybase\_Database\_Summary.
- DBD15 will exclude cursor execution for dataset Sybase\_Database\_Detail.
- KOYSEGD will exclude cursor execution for dataset Sybase\_Segment\_Detail.

### What to do next

When the configuration is completed, you can start the monitoring agent and connect the monitoring agent to the Performance Management server.

To start the Sybase agent, use the agent command `sybase-agent` command.

To connect the Sybase agent to the Performance Management server, use the `agent2server` command.

## Configuring the agent by using the silent response file

You can configure the Monitoring Agent for Sybase Server by using the silent response file.

### Before you begin

The Sybase agent does not support remote configuration. Hence, you need to ensure that the Sybase Server is installed on the same host where the Sybase agent is installed.

The Sybase agent supports Sybase Server version 15.7 and 16.0 only.

The user ID that is used to connect to the database server is created.

### About this task

The Sybase agent is a multiple instance agent, you must configure and start each agent instance manually.

You must edit the silent response file and run the agent command to configure the monitoring agent.

### Procedure

1. Edit the silent response file.

- **Windows**

Silent response file is at: `install_dir\samples\sybase_silent_config.txt`.

- **UNIX**

Silent response file is at: `install_dir/samples/sybase_silent_config.txt`.

Where `install_dir` is the agent installation directory.

a) For the `Home Directory` parameter, specify the path of the Sybase Server home directory.

- **Windows**

The example of `Home Directory` is `\opt\sybase`.

- **UNIX**

The example of `Home Directory` is `/opt/sybase`.

b) For the `ASE Directory` parameter, specify the path of the database server ASE.

- **Windows**

The example of `ASE Directory` is `\opt\sybase\ASE-12_5`.

- **UNIX**

The example of `ASE Directory` is `/opt/sybase/ASE-12_5`.

c) For the `Open Client Directory` parameter, specify the Sybase open client installation location.

- **Windows**

The example of Open Client Directory is `\opt\sap\ocs-16_0`.

- **UNIX**

The example of Open Client Directory is `/opt/sap/ocs-16_0`.

- d) For the `USER ID` parameter, specify the user ID that is used by the monitoring agent to connect to the Sybase Server.

The default `USER ID` is `tivoli`.

- e) For the `PASSWORD` parameter, specify the password of the user ID that is used by the monitoring agent to connect to the Sybase Server.

- f) For the `VERSION` parameter, specify the Sybase Server version.

The Sybase agent supports Sybase Server versions 15.7 and 16.0 only.

- g) For the `ERROR LOG FILE` parameter, specify the fully qualified file name of the error log file for the Sybase Server.

- **Windows**

The example of the `ERROR LOG FILE` is `\opt\sap\ASE-16_0\install\servername.log`.

- **UNIX**

The example of the `ERROR LOG FILE` is `/opt/sap/ASE-16_0/install/servername.log`.

Where `servername` is the Sybase Server name.

- h) For the `EXTENDED` parameter, specify the extended parameter that is used by support to exclude certain cursor execution. Optionally, leave it blank to execute all cursors.

The options for the `EXTENDED` are `DBD2`, `DBD15`, `KOYSEGD`.

- `DBD2` will exclude cursor execution for datasets `Sybase_Database_Detail` and `Sybase_Database_Summary`.
- `DBD15` will exclude cursor execution for dataset `Sybase_Database_Detail`.
- `KOYSEGD` will exclude cursor execution for dataset `Sybase_Segment_Detail`.

2. Save the silent response file.

3. Run the following agent command to configure the monitoring agent.

- **Windows**

```
install_dir\bin\sybase-agent.bat config instance_name
install_dir\samples\sybase_silent_config.txt
```

- **UNIX**

```
install_dir/bin/sybase-agent.sh config instance_name
install_dir/samples/sybase_silent_config.txt
```

Where:

- `install_dir` is the agent installation directory.
- `instance_name` is the Sybase Server name.

## What to do next

When the configuration is completed, you can start the monitoring agent and connect the monitoring agent to the Performance Management server.

To start the Sybase agent, use the **sybase-agent** command.

To connect the Sybase agent to the Performance Management server, use the **agent2server** command.

## Disabling dirty reads for query

The Sybase agent enables `dirty reads` for its query execution by default to prevent locking. The variable `COLL_USE_NOLOCK` is used to enable or disable the query `dirty reads`. When `dirty reads` is enabled, query is executed with isolation level zero to avoid locking. If you wish to disable the `dirty reads` for agent query, you can set the variable `COLL_USE_NOLOCK` to zero.

### Before you begin

To disable `dirty reads` for agent query, ensure that the agent is installed.

### About this task

The Sybase agent enables `dirty reads` by default. To disable the `dirty reads` for the agent query, complete the following steps.

### Procedure

1. Stop the agent.
2. Set the variable `COLL_USE_NOLOCK` to zero.
  - **UNIX**
    - a. Add **`COLL_USE_NOLOCK=0`** in `CANDLEHOME/config/.oy.environment` file.
    - b. Save and close the file.
  - **Windows**
    - a. Locate the agent instance file `CANDLEHOME\TMAITM6_x64\KOYENV_INSTANCENAME`.
    - b. Add the following line in the file.  
**`COLL_USE_NOLOCK=0`**
    - c. Save and close the file.
3. Start the agent.

The `CANDLEHOME` is the agent installation directory.  
The `INSTANCENAME` is the agent instance name.

## Configuring Tomcat Monitoring

You can configure the Monitoring Agent for Tomcat with the default or custom settings to monitor the resources of Tomcat application servers. The agent can be configured on Windows and Linux systems.

### Before you begin

- Enable JMX remote for the monitored Tomcat server. Set the port. For instructions, see [https://tomcat.apache.org/tomcat-6.0-doc/monitoring.html#Enabling\\_JMX\\_Remote](https://tomcat.apache.org/tomcat-6.0-doc/monitoring.html#Enabling_JMX_Remote).
- Ensure that the Tomcat server that you want to monitor is up and running.

The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history” on page 52](#).

### About this task

The Tomcat agent is a multiple instance agent; you must create the first instance and start the agent manually. The managed system name includes the instance name that you specify, for example, `instance_name: host_name: pc`, where `pc` is your two character product code. The managed system name is limited to 32 characters. The instance name that you specify is limited to 28 characters that excludes the length of your host name. For example, if you specify `TOMCAT2` as your instance name, your

managed system name is TOMCAT2:hostname:OT. If you specify a long instance name, the managed system name is truncated, and the agent code is not displayed completely.

To avoid permission issues when you configure the agent, be sure to use the same root user or non-root user ID that was used for installing the agent. If you installed your agent as a selected user and want to configure the agent as a different user, see [“Configuring agents as a non-root user” on page 231](#). If you installed and configured your agent as a selected user and want to start the agent as a different user, see [“Starting agents as a non-root user” on page 230](#).

## Configuring Tomcat agent with the default settings

You can use the default settings of the Tomcat agent to monitor the Tomcat server. You do not need to provide any configuration information other than the new instance name.

### Before you begin

Before you configure the agent with the default settings, ensure that the following prerequisites are met:

- The agent is installed in the default directory.
- The JMX service URL uses the 8686 port.
- The Tomcat server is configured without the JMX authorization.

### About this task

**Remember:** When you configure the agent with the default settings, the collection of transaction tracking and deep-dive diagnostics data is not enabled.

### Procedure

1. Run the following command:

```
Linux install_dir/bin/tomcat-agent.sh config instance_name install_dir/samples/tomcat_silent_config.txt
```

```
Windows install_dir/bin/tomcat-agent.bat config instance_name install_dir/samples/tomcat_silent_config.txt
```

Where

#### ***install\_dir***

The installation directory of the Tomcat agent.

#### ***instance\_name***

The name that you want to give to the instance.

2. Run the following command to start the agent:

```
Linux install_dir/bin/tomcat-agent.sh start instance_name
```

```
Windows install_dir/bin/tomcat-agent.bat start instance_name
```

### What to do next

Log in to the Cloud App Management console to view data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring Tomcat agent on Windows systems

You can configure the agent on Windows operating systems by using the **IBM Performance Management** window.

### Before you begin

Ensure that the following prerequisites are met:

- Java is installed on the Tomcat Server where the agent is installed.
- JMX Remote is enabled for the Tomcat Server. For details, see [Enabling JMX Remote](#).

- The Tomcat Server is up and running.

### About this task

This topic explains configuring the agent by using the agent configuration panel.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Monitoring Agent for Tomcat**.
3. Click **Configure agent**.



**Attention:** If **Configure agent** is unavailable, click **Reconfigure**.

4. In the **Instance Name** window, specify a unique name for the Tomcat agent instance, and click **OK**.

**Restriction:** The MSN must not exceed 32 characters.

5. In the **SERVER NAME** field, enter a unique name to identify the Tomcat Server that is being monitored.
6. In the **Java Parameter Settings** window, complete one of the following steps:
  - Click **Next** to accept the default location where Java is installed. The default installation path is C:\IBM\APM\java\java80\_x64\jre.
  - In the **Java Home** field, specify the path when IBM Java is installed at a different path.
7. In the **JSR-160-Complaint Server** window, specify the details of the following parameters:
  - a) In the **JMX user ID** field, specify the ID of the user that is used to connect to the Tomcat MBean server when the JMX authorization is enabled in Tomcat.
  - b) In the **JMX password** field, specify the password of the JMX user when the JMX authorization is enabled in Tomcat.
  - c) In the **JMX service URL** field, enter the URL that is used for connecting to the Tomcat MBean server.  
The format of the URL is `service:jmx:rmi:///jndi/rmi://host_name:port_number/jmxrmi`. The default URL is valid when the server runs on the local host and uses the 8686 port as a JMX port. You can modify the host name and port number in the URL, keeping the same format.
  - d) From the **Data Collector Configuration** list, select Yes if you want to enable collection of transaction tracking and deep dive data.
8. In the **Monitoring Agent for Tomcat** window, right-click the Tomcat agent instance, and click **Start**.
9. Enable the collection of Transaction Tracking and Deep Dive data and restart the Tomcat Server.

### What to do next

Log in to the Cloud App Management console to view data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

## Configuring Tomcat agent on Linux systems

You run the configuration script and respond to prompts to configure the Tomcat agent on Linux systems.

### Before you begin

- JMX Remote is enabled for the Tomcat Server. For details, see [Enabling JMX Remote](#).
- The Tomcat Server is up and running.

### Procedure

1. Run the following command:  
`install_dir/bin/tomcat-agent.sh config instance_name`  
 Where `instance_name` is the name that you want to give to the instance.

2. When you are prompted to specify a value for `SERVER`, specify a unique name to identify the Tomcat Server that is being monitored, and press Enter.
3. When you are prompted to specify a value for `Java home`, press Enter to accept the default location where the Java virtual machine is installed. The default location is `/opt/ibm/apm/agent/JRE/1x8266/jre`. If the agent is not installed in the default directory, specify `install_dir/JRE/1x8266/jre`.
4. When you are prompted to specify a value for `JMX user ID`, specify the ID of the user who connects to the Tomcat MBean server. If the JMX authorization is not enabled, press Enter.
5. When you are prompted to specify a value for `JMX password`, specify the password of the JMX user and confirm it. If JMX authorization is not enabled, press Enter.
6. When you are prompted to specify a value for `JMX service URL`, press Enter to accept the default URL or specify another service URL for connecting to the Tomcat MBean server. The format of the URL is `service:jmx:rmi:///jndi/rmi://host_name:port_number/jmxrmi`. The default URL is valid when the server runs on the local host and uses the 8686 port as a JMX port. You can modify the host name and the port in the URL, keeping the same format.
7. When you are prompted to specify a value for `Data Collector Configuration`, specify `1` and press Enter to enable collection of transaction tracking and deep dive data.
8. Run the following command to start the agent:  
`install_dir/bin/tomcat-agent.sh start instance_name`
9. Enable the collection of Transaction Tracking and Deep Dive data, restart the Tomcat Server.

### What to do next

Log in to the Cloud App Management console to view data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring Tomcat agent using silent response file

You can use the silent response file to configure the Tomcat agent without responding to prompts.

### Procedure

1. In a text editor, open the `tomcat_silent_config.txt` file that is available at the following path:  
`install_dir/samples`
2. For the **KOT\_SERVER** parameter, specify a unique name to identify the Tomcat Server that is being monitored.
3. For the **Java home** parameter, specify the path where the Java virtual machine is installed. The default location is `/opt/ibm/apm/agent/JRE/1x8266/jre`. If the agent is not installed in the default directory, specify `install_dir/JRE/1x8266/jre`.
4. For the **JMX user ID** parameter, specify the ID of the user that is used to connect to the Tomcat MBean server. You must specify a value for this parameter when the JMX authorization is enabled in Tomcat.
5. For the **JMX password** parameter, specify the password of the JMX user. You must specify a value for this parameter when the JMX authorization is enabled in Tomcat.
6. For the **JMX service URL** parameter, specify the service URL for connecting to the Tomcat MBean server. The format of the URL is `service:jmx:rmi:///jndi/rmi://host_name:port_number/jmxrmi`. The default URL is valid when the server runs on the local host and uses the 8686 port as a JMX port. You can modify the host name and the port number in the URL, keeping the same format.
7. For the **KOT\_DCCONFIGURATION** parameter, specify `Yes` if you want to enable collection of transaction tracking and deep dive data.
8. Save and close the `tomcat_silent_config.txt` file, and run the following command to update the agent configuration settings:

```
Linux install_dir/bin/tomcat-agent.sh config instance_name install_dir/samples/tomcat_silent_config.txt
```

```
Windows install_dir/bin/tomcat-agent.bat config instance_name install_dir/samples/tomcat_silent_config.txt
```

Where *instance\_name* is the name that you want to give to the instance, and *install\_dir* is the installation directory of the Tomcat agent.

9. Run the following command to start the agent:

```
Linux install_dir/bin/tomcat-agent.sh start instance_name
```

```
Windows install_dir/bin/tomcat-agent.bat start instance_name
```

10. If you enable the collection of transaction tracking and deep dive data, restart the Tomcat Server.

### What to do next

Log in to the Cloud App Management console to view data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring VMware VI monitoring

After installing the Monitoring Agent for VMware VI, you must create the first instance, and manually start the agent so that the agent can collect data of the VMware Virtual Infrastructure that is being monitored.

### Before you begin

- Review the hardware and software prerequisites.
- Create a user ID in your VMware Virtual Infrastructure. The agent uses this user ID to connect to the VMware vCenter for monitoring the VMware Virtual Infrastructure. Ensure that you have the "System.View" and "System.Read" privileges on all the vCenters and ESX servers that are being monitored. For information about how to create the user ID, see the VMware documentation for managing users, groups, permissions, and roles.
- Determine whether the vCenter is configured for SSL communication. If it is configured, then you must configure the VMware VI agent to use SSL for communicating with the vCenter.
  - To determine whether the vCenter uses SSL for communication, use the `https://vCenterIPAddress` URL to access the vCenter. If you can access the vCenter, then it indicates that the vCenter uses SSL to communicate over the network.
  - To configure the VMware VI agent to use SSL for communicating with the vCenter, complete the steps that are described in [“Enabling SSL communication with VMware VI data sources”](#) on page 516.
- Decide the number of agent instances that you need to monitor your VMware Virtual Infrastructure. For information about sizing the agent instances according to your monitoring environment, see [“Sizing and planning the VMware VI agent deployment”](#) on page 516.

### About this task

The VMware VI agent is a multiple instance agent. Unlike a single instance agent, for which you can configure the agent to monitor and collect data for only one monitored application, the VMware VI agent can have multiple configured instances that connect to multiple vCenter servers and remotely monitor your VMware Virtual Infrastructure.

The configuration parameters define the VMware VI data sources that are monitored and define a connection to either the VMware vCenter, vCenter Server Appliance, or to an individual VMware ESX server. To know the supported versions of these applications, see the [Software Product Compatibility Reports for the VMware VI agent](#).

You must manually configure the agent to view data for all the agent attributes.

- To configure the agent on Windows operating systems, you can use the **IBM Performance Management** window or the silent response file.
- To configure the agent on Linux operating systems, you can run the script and respond to prompts, or use the silent response file.

## Sizing and planning the VMware VI agent deployment

The number of agent instances that you can configure on a single system depends on the availability and utilization of resources on the system.

The following table categorizes the VMware environment into various sizes with the required Java heap size:

| VMware environment size       | Number of ESX servers                                                                  | Java heap size               |
|-------------------------------|----------------------------------------------------------------------------------------|------------------------------|
| <b>Small environment</b>      | A vCenter server that manages up to 125 ESX(i) servers and 300 - 1500 guests.          | <b>-Xmx2048m</b><br>(2 GB)   |
| <b>Medium environment</b>     | A vCenter server that manages between 125 - 250 ESX(i) servers and 1500 - 4000 guests. | <b>-Xmx4096m</b><br>(4 GB)   |
| <b>Large environment</b>      | A vCenter server that manages between 250 - 500 ESX(i) servers and 4000 - 7500 guests. | <b>-Xmx8192m</b><br>(8 GB)   |
| <b>Very large environment</b> | A vCenter server that manages more than 500 ESX(i) servers and more than 7500 guests.  | <b>-Xmx16384m</b><br>(16 GB) |

To increase the heap size for the Java data provider, complete the steps that are described in [“Increasing the Java heap size”](#) on page 522.

For the agent instances to successfully monitor the environment, the server on which you install the agent, must have adequate memory resources to accommodate the data that is collected by these agent instances. A single instance of the VMware VI agent requires approximately 300 - 400 MB to monitor a small environment. See the following guidelines about the number of agent instances to be configured:

- Use a single instance to monitor a single vCenter. Do not use the same instance to monitor multiple vCenters.
- In a non-cluster environment, use a single instance to monitor a maximum of 8 small ESX servers (100 - 200 virtual machines in one ESX server). Do not configure multiple individual ESX servers under the single agent instance.
- Use multiple agent instances of the VMware VI agent to monitor an environment that contains multiple vCenters. Before you configure multiple instances, ensure that you have adequate memory resources on the system where you install the agent.

## Enabling SSL communication with VMware VI data sources

Before you configure the agent to securely communicate with its VMware VI data sources by using SSL, you must add a data source SSL certificate to the certificate truststore of the agent.

### About this task

**Important:** The following information applies only if the agent is configured to validate SSL certificates.

If the SSL certificate validation is turned off, the VMware VI agent connects to VMware data sources even if their SSL certificates are expired, untrusted, or invalid. However, turning off SSL certificate validation is potentially not secure and must be done with care.



If a VMware data source uses an SSL certificate that is signed by a common Certificate Authority (for example, Verisign, Entrust, or Thawte), then it is not necessary to add certificates to the VMware VI agent certificate truststore. However, if the data source uses a certificate that is not signed by a common Certificate Authority, as is the case by default, you must add the certificate to the truststore to allow the agent to successfully connect and collect data.

**Note:**

1. The default VMware certificate file is named `rui.crt`.
2. For a Virtual Center, the SSL certificate file is located by default in the following path:  
`C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\SSL`
3. For an ESX server, the SSL certificate file is located by default in the `/etc/vmware/ssl` directory.

**Procedure**

1. Copy the certificate file from your data source to the agent computer.
2. On the agent computer, place the certificate file in a directory of your choice. Do not overwrite the certificate files. Use a unique file name and a label for each certificate that you add.
3. Use the `keytool` command to add the data source certificate to the certificate truststore of the agent:

```
keytool -import -noprompt -trustcacerts -alias CertificateAlias -file \  
CertificateFile -keystore Truststore -storepass TruststorePassword
```

Where

**CertificateAlias**

Unique reference for each certificate added to the certificate truststore of the agent, for example, an appropriate alias for the certificate from `datasource.example.com` is `datasource`.

**CertificateFile**

Complete path and file name to the VMware data source certificate to add to the truststore.

**Truststore**

Complete path and file name to the VMware VI agent certificate database. Use the following path and file name:

- **Windows** (64 bit): `install_dir\tmaitm6_x64\kvm.truststore`
- **Linux** (64 bit): `install_dir/1x8266/vm/etc/kvm.truststore`

**TruststorePassword**

ITMVMWAREVI is the default password for the VMware VI agent truststore. To change this password, consult the Java Runtime documentation for information about the tools to use.

**Important:** To use the `keytool` command, the Java Runtime bin directory must be in your path. Use the following commands:

- **Windows** (64 bit): `set PATH=%PATH%;install_dir\java\java70_x64\jre\bin`
- **Linux** (64 bit): `PATH="$PATH" : /opt/ibm/apm/agent/JRE/1x8266/bin`

4. After you add all the data source certificates, start the monitoring agent.

**What to do next**

Complete the agent configuration.

## Configuring the agent on Windows systems

You can configure the agent on Windows operating systems by using the **IBM Performance Management** window. After you update the configuration values, you must start the agent to save the updated values.

### About this task

The VMware VI agent provides default values for some parameters. You can specify different values for these parameters.

### Procedure

1. Click **Start > All Programs > IBM Monitoring agents > IBM Performance Management**.
2. In the **IBM Performance Management** window, right-click **Monitoring Agent for VMware VI**, and then click **Configure agent**.

**Remember:** After you configure the agent for the first time, the **Configure agent** option is disabled. To configure the agent again, click **Reconfigure**.

3. In the Monitoring Agent for VMware VI window, complete the following steps:
  - a) Enter a unique name for the VMware VI agent instance, and click **OK**.
  - b) On the **Data Provider** tab, specify values for the configuration parameters, and then click **Next**.
  - c) On the **Data Source** tab, specify values for the configuration parameters, and then click **Next**.

The VMware VI agent is a multi-data source agent. You can monitor multiple data sources from the same agent.

- If you want to configure a new data source, click **New**.
- If you want to delete an existing data source, click **Delete**.

For information about the configuration parameters in each tab of the VMware VI agent window, see the following topics:

- [Configuration parameters for the data provider](#)
- [Configuration parameters for the data source](#)

4. In the **IBM Performance Management** window, right-click the instance that you configured, and then click **Start**.

### What to do next

- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI”](#) on page 176.

If you need help with troubleshooting, see the [IBM Cloud APM Forum](#) on developerWorks.

- If you are monitoring a large VMware environment with more than 500 ESX hosts, you might need to increase the heap size for the Java data provider. For more information, see [“Increasing the Java heap size”](#) on page 522.

## Configuring the agent by responding to prompts

To configure the agent on Linux operating systems, you must run the script and respond to prompts.

### Procedure

- To configure the agent by running the script and responding to prompts, complete the following steps:
  - a) On the command line, run the following command:

```
install_dir/bin/vmware_vi-agent.sh config instance_name
```

```
Example /opt/ibm/apm/agent/bin/vmware_vi-agent.sh config instance_name
```

Where

**instance\_name**

Name that you want to give to the instance.

**install\_dir**

Path where the agent is installed.

b) Respond to the prompts by referring to the following topics:

- [“Configuration parameters for the data provider” on page 521](#)
- [“Configuration parameters for the data source” on page 520](#)

c) Run the following command to start the agent:

```
install_dir/bin/vmware_vi-agent.sh start instance_name
```

Example **/opt/ibm/apm/agent/bin/vmware\_vi-agent.sh start instance\_name**

**What to do next**

- Log in to the Cloud App Management console to view the data that is collected by the agent in the dashboards. For information about using the Cloud App Management console, see [“Starting the Cloud App Management UI” on page 176](#).

If you need help with troubleshooting, see the [IBM Cloud APM Forum](#) on developerWorks.

- If you are monitoring a large VMware environment with more than 500 ESX hosts, you might need to increase the heap size for the Java™ data provider. For more information, see [“Increasing the Java heap size” on page 522](#).

**Configuring the agent by using the silent response file**

The silent response file contains the agent configuration parameters. You can edit the silent response file to modify the values of these parameters, and run the configuration script to create an instance and update the agent configuration values. This mode of configuration is called the silent mode.

**About this task**

The silent response file contains the agent configuration parameters with default values that are defined for some parameters. You can edit the silent response file to specify different values for the configuration parameters.

After you update the configuration values in the silent response file, you must run the configuration script to configure the agent with these updated values.

**Procedure**

- To configure the VMware VI agent in the silent mode, complete the following steps:
  - a) In a text editor, open the `vmware_vi_silent_config.txt` file that is available at the following path:

```
– Linux install_dir/samples/vmware_vi_silent_config.txt
```

```
Example /opt/ibm/apm/agent/samples/vmware_vi_silent_config.txt
```

```
– Windows install_dir\samples\vmware_vi_silent_config.txt
```

```
Example C:\IBM\APM\samples\vmware_vi_silent_config.txt
```

- b) In the `vmware_vi_silent_config.txt` file, specify values for all mandatory parameters. You can also modify the default values of other parameters.

For information about the configuration parameters, see the following topics:

- [“Configuration parameters for the data provider” on page 521](#)
- [“Configuration parameters for the data source” on page 520](#)

c) Save and close the `vmware_vi_silent_config.txt` file, and run the following command:

– **Linux** `install_dir/bin/vmware_vi-agent.sh config instance_name install_dir/samples/vmware_vi_silent_config.txt`

Example `/opt/ibm/apm/agent/bin/vmware_vi-agent.sh config instance_name /opt/ibm/apm/agent/samples/vmware_vi_silent_config.txt`

– **Windows** `install_dir\bin\vmware_vi-agent.bat config instance_name install_dir\samples\vmware_vi_silent_config.txt`

Example `C:\IBM\APM\bin\ vmware_vi-agent.bat config instance_name C:\IBM\APM\samples\vmware_vi_silent_config.txt`

Where

**instance\_name**

Name that you want to give to the instance.

**install\_dir**

Path where the agent is installed.

**Important:** Ensure that you include the absolute path to the silent response file. Otherwise, the agent data is not shown in the dashboards.

d) Run the following command to start the agent:

– **Linux** `install_dir/bin/vmware_vi-agent.sh start instance_name`

Example `/opt/ibm/apm/agent/bin/vmware_vi-agent.sh start instance_name`

– **Windows** `install_dir\bin\vmware_vi-agent.bat start instance_name`

Example `C:\IBM\APM\bin\vmware_vi-agent.bat start instance_name`

### What to do next

- Log in to the Cloud APM console to view the data that is collected by the agent in the dashboards. For information about using the Cloud APM console, see [“Starting the Cloud App Management UI”](#) on page 176.

If you need help with troubleshooting, see the [IBM Cloud APM Forum](#) on developerWorks.

- If you are monitoring a large VMware environment with more than 500 ESX hosts, you might need to increase the heap size for the Java™ data provider. For more information, see [“Increasing the Java heap size”](#) on page 522.

### Configuration parameters for the data source

When you configure the VMware VI agent, you can change the default values of the parameters for the data source, such as the address, user id, and password of the data source.

The following table contains detailed descriptions of the configuration parameters for the data source.

| Parameter name | Description                | Mandatory field |
|----------------|----------------------------|-----------------|
| Data Source ID | The ID of the data source. | Yes             |

Table 73. Names and descriptions of the configuration parameters for the data source (continued)

| Parameter name                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Mandatory field |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Data Source Address               | <p>Address of the data source.</p> <p>If you do not want the agent to validate the SSL certificates, set the value to the host name or IP address of the VMware Virtual Center or ESX server that is being monitored.</p> <p>If you want the agent to validate the SSL certificates when using SSL to communicate over the network, configure the agent by using the Subject Alternative Name that is provided in the certificate.</p> <p>To view the subject alternative name of the data center, complete the following steps:</p> <ol style="list-style-type: none"> <li>1. Open the certificate.</li> <li>2. In the <b>Certificate</b> window, click the <b>Details</b> tab.</li> <li>3. Select <b>Subject Alternative Name</b>, and use the value of DNS Name. For example, if the value of DNS Name is "ibmesx3v3vc.ITMfVS.com", then use the "ibmesx3v3vc.ITMfVS.com" value for the host name.</li> </ol> | Yes             |
| Use SSL Connection to Data Source | <p>Indicates whether the agent uses an SSL connection to connect to the data sources of the VMware Virtual Infrastructure.</p> <p>Specify Yes if the agent uses an SSL connection to connect to the data sources. Otherwise, specify No. The default value is Yes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Yes             |
| Data Source User ID               | The user ID that has sufficient privileges to collect monitoring data, and is known to the data source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Yes             |
| Data Source Password              | The password of the user ID that is configured for accessing the data source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Yes             |
| Confirm Data Source Password      | The same password that you specified in the <b>Data Source Password</b> field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                 |

### Configuration parameters for the data provider

When you configure the VMware VI agent, you can change the default values of the parameters for the data provider, such as the maximum number of data provider log files, the maximum size of the log file, and the level of detail that is included in the log file.

The following table contains detailed descriptions of the configuration parameters for the data provider.

Table 74. Names and descriptions of the configuration parameters for the data provider

| Parameter name | Description                                                                                                                                                                                                                                                                          | Mandatory field |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Instance Name  | <p>The name of the instance.</p> <p><b>Restriction:</b> The <b>Instance Name</b> field displays the name of the instance that you specify when you configure the agent for the first time. When you configure the agent again, you cannot change the instance name of the agent.</p> | Yes             |

| Parameter name                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Mandatory field |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Valid SSL Certificates                       | <p>Indicates whether the agent validates SSL certificates when the agent uses SSL to communicate over the network.</p> <p>Set the value to Yes if you want the agent to validate SSL certificates when the agent uses SSL to communicate over the network. Set the value to No to prevent the agent from validating SSL certificates. The default value is Yes.</p> <p>For information about adding a data source SSL certificate to the certificate truststore of the agent, see <a href="#">“Enabling SSL communication with VMware VI data sources”</a> on page 516.</p> | Yes             |
| Maximum number of Data Provider Log Files    | The maximum number of log files that the data provider creates before it overwrites the previous log files. The default value is 10.                                                                                                                                                                                                                                                                                                                                                                                                                                        | Yes             |
| Maximum Size in KB of Each Data Provider Log | The maximum size in KB that a data provider log file must reach before the data provider creates a new log file. The default value is 5190 KB.                                                                                                                                                                                                                                                                                                                                                                                                                              | Yes             |
| Level of Detail in Data Provider Log         | The level of detail that can be included in the log file that the data provider creates. The default value is INFO. The following values are valid: OFF, SEVERE, WARNING, INFO, FINE, FINER, FINEST, and ALL.                                                                                                                                                                                                                                                                                                                                                               | Yes             |

## Increasing the Java heap size

After you configure the VMware VI agent, if you are monitoring a large VMware Virtual Infrastructure environment, then you might need to increase the heap size for the Java™ data provider.

### About this task

The default maximum heap size for the Java data provider is 256 megabytes. You must set the maximum heap size to an appropriate value that depends on the size of the VMware environment. For information about the heap sizes that are required for the various VMware environments, see [“Sizing and planning the VMware VI agent deployment”](#) on page 516.

**Important:** The system, on which you install and configure the VMware VI agent, must have adequate memory space to accommodate the required heap size.

If any of the following problems arise, then you might need to increase the heap size:

- The Java data provider stops because of a javacore problem, and creates a file that is named `javacore.date.time.number.txt` in the `CANDLEHOME\tmaitm6_x64` directory.
- The `javacore.date.time.number.txt` file contains the string `java/lang/OutOfMemoryError`.

### Procedure

#### Windows

Complete the following steps to set a value of 1 GB as the heap size:

1. Open the `%CANDLE_HOME%\TMAITM6_x64\kvm_data_provider.bat` file.
2. Add the following line before the line that starts with `KVM_JVM_ARGS=%KVM_CUSTOM_JVM_ARGS...`:

```
SET KVM_CUSTOM_JVM_ARGS=-Xmx1024m
```

3. Restart the agent.

- **Linux**

Complete the following steps to set a value of 1 GB as heap size:

1. Open the `$CANDLEHOME/1x8266/vm/bin/kvm_data_provider.sh` file.
2. Add the following line before the line that starts with `KVM_JVM_ARGS="$KVM_CUSTOM_JVM_ARGS...`:

```
KVM_CUSTOM_JVM_ARGS=-Xmx1024m
```

3. Restart the agent.

## Configuring WebSphere Applications monitoring

The WebSphere Applications agent does not need any configuration after agent installation, unless you want to change the default port. However, you must configure the data collector, which is a component of the agent, to set up monitoring for your WebSphere environment.

### Before you begin

The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history” on page 52](#).

### About this task

**Remember:** Diagnostics and transaction tracking data are not yet supported by Cloud App Management. Do NOT enable diagnostics data or transaction tracking for the data collector.

### Procedure

- (Fast track) To quickly set up the environment for monitoring, see [“Fast track: Configuring the data collector for WebSphere Applications agent” on page 523](#) for a simplified configuration flow.
- (Simple configuration) For a complete configuration flow for a pure Cloud App Management environment, see [“Configuring the data collector with the simple configuration utility” on page 526](#).
- (Full configuration) To configure the data collector with more customization options, use the full configuration utilities. For instructions, see [“Configuring or reconfiguring the data collector with full configuration utilities” on page 527](#).
- (Silent configuration) To deploy the same monitoring for many application server instances, configure the data collector in silent mode. For instructions, see [“Configuring the data collector in silent mode” on page 534](#).
- (WebSphere Portal Server) To monitor WebSphere Portal Server instances, use the advanced configuration procedure. For instructions, see [“Configuring or reconfiguring the data collector with full configuration utilities” on page 527](#).
- (Manual configuration) If you cannot use the provided configuration utilities to configure the data collector for WebSphere Applications agent, manually configure the data collector in the WebSphere Administrative Console. For instructions, see [“Manually configure the data collector if the configuration utilities fail” on page 542](#).

### Fast track: Configuring the data collector for WebSphere Applications agent

The WebSphere Applications agent does not need any configuration after agent installation. However, you must configure the data collector, which is a component of the agent, to set up monitoring for your WebSphere environment.

### Before you begin

1. Install the WebSphere Applications agent on the system where the application server to be monitored is installed and running.

## 2. Check the user access requirements.

- **Windows** Use the administrator ID that is used to install the application server to configure the data collector. Make sure that this user ID has full write permission the data collector home directory, `install_dir\dchome\7.3.0.14.09`.
- **Linux** **UNIX** Use the user ID that is used to install the application server to configure the data collector. Make sure that this user ID has read and write permissions to the following sub-directories within `install_dir/yndchome/7.3.0.14.09`:
  - bin
  - data
  - runtime

### About this task

A simple configuration utility, `simpleconfig`, is used in this procedure to provide the basic configuration of data collector.

The `simpleconfig` utility configures the data collector with default settings. To configure the data collector with more customization options, use the full configuration utility, `config`, in the same directory. For instructions, see [“Configuring or reconfiguring the data collector with full configuration utilities”](#) on page 527.

In most cases, the `simpleconfig` utility is sufficient. For more complex environment, you can use the `config` configuration utility to configure the data collector. If the `simpleconfig` utility fails, use the `config` utility instead.

### Procedure

1. Log in to the system with the user ID that is used to install the application server.
2. Change to the `bin` directory within the data collector home directory.

- **Windows** `install_dir\dchome\7.3.0.14.09\bin`
- **Linux** **UNIX** `install_dir/yndchome/7.3.0.14.09/bin`

3. Run the following simple configuration utility:

- **Windows** `simpleconfig.bat`
- **Linux** **UNIX** `./simpleconfig.sh`

4. Follow the prompts to continue with the data collector configuration.

You are required to do some or all of the following things, depending on the application server settings:

- For traditional WebSphere Application Server:
    - Select the auto-discovered WebSphere installation directory or manually specify the installation directory.
    - Select the WebSphere Application Server profile to monitor.
    - Select the security properties profile to use or provide the user name and password of the WebSphere administrative console (if security is enabled for the application server).
  - For WebSphere Application Server Liberty:
    - Specify the full path of the Liberty home directory that contains `bin` and `servers` directories. For example, `/opt/ibm/wlp`.
    - Specify the home directory of the JRE that is used by Liberty.
5. After the data collector configuration completes, restart the application server.
    - a) Go to the `bin` directory under the home directory for the application server profile. For example, `opt/IBM/WebSphere/AppServer/profiles/profile_name/bin`.



b) Stop the application server by entering the **stopServer** command in the command console.

- **Linux** | **UNIX** `./stopServer.sh server_name`
- **Windows** `stopServer.bat server_name`

c) When prompted, enter the user ID and password of WebSphere administrative console administrator.

d) Start the application server again by entering the **startServer** command in the command console.

- **Linux** | **UNIX** `./startServer.sh server_name`
- **Windows** `startServer.bat server_name`

## Results

The data collector is configured to monitor the application server instance.

Now, you can log in to the Cloud App Management console to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).

## Checking user access requirements

The WebSphere Applications agent has some user access requirements for the user ID that is to configure the data collector.

### About this task

Use the ID that is used to install the application server to configure the data collector after you grant appropriate permissions for the application server installation ID.

### Procedure

- **Windows** Use the administrator ID that is used to install the application server to configure the data collector. Make sure that this user ID has full write permission the data collector home directory, `install_dir\dchome\7.3.0.14.09`.
- **Linux** | **UNIX** Use the user ID that is used to install the application server to configure the data collector. Make sure that this user ID has read and write permissions to the following sub-directories within `install_dir/yndchome/7.3.0.14.09`:
  - bin
  - data
  - logs
  - runtime

**Remember:** If you use different user IDs to install application servers, you might need to use different user IDs to configure the data collector. After you configure the data collector for the first time, grant the write permission to the following files every time you use a different user ID to configure the data collector, where *profile\_name* is the application server profile name:

- `install_dir/yndchome/7.3.0.14.09/data/findservers.inputlist`
- `install_dir/yndchome/7.3.0.14.09/data/profile_name.findservers.progress`
- `install_dir/yndchome/7.3.0.14.09/data/config.inputlist`
- `install_dir/yndchome/7.3.0.14.09/runtime/custom/connections.properties`

## Configuring the data collector with the simple configuration utility

The WebSphere Applications agent starts automatically after installation, but you must manually configure the data collector, which is a component of the agent, to monitor application server instances.

### Before you begin

Make sure that the user access requirements are met in your environment. For instructions, see [“Checking user access requirements”](#) on page 525.

### About this task

For the WebSphere Applications agent, the `dc_home` variables refer to the home directory of the data collector. The location of the `dc_home` variable on each operating system is as follows:

- **Windows** `install_dir\dchome\7.3.0.14.09`
- **Linux** | **UNIX** `install_dir/yndchome/7.3.0.14.09`

### Procedure

1. Log in to the system with the user ID that is used to install the application server.
2. Change to the bin directory within the data collector home directory.

- **Windows** `install_dir\dchome\7.3.0.14.09\bin`
- **Linux** | **UNIX** `install_dir/yndchome/7.3.0.14.09/bin`

3. Run the following simple configuration utility:

- **Windows** `simpleconfig.bat`
- **Linux** | **UNIX** `./simpleconfig.sh`

The **simpleconfig** utility automatically discovers the home directories of the application servers.

4. Follow the prompts to continue with the data collector configuration.

You are required to do the following things, depending on the application server settings:

- For traditional WebSphere Application Server:
    - Select the auto-discovered WebSphere installation directory or manually specify the installation directory.
    - Select the WebSphere Application Server profile to monitor.
    - Select the security properties profile to use or provide the user name and password of the WebSphere administrative console (if security is enabled for the application server).
  - For WebSphere Application Server Liberty:
    - Specify the full path of the Liberty home directory that contains the `bin` and `servers` directories (for example, `/opt/ibm/wlp`).
    - Specify the home directory of the JRE that is used by Liberty.
5. If possible, restart the application server instance after the data collector configuration completes.
    - a) Go to the `bin` directory under the home directory for the application server profile. For example, `opt/IBM/WebSphere/AppServer/profiles/profile_name/bin`.
    - b) Stop the application server by entering the **stopServer** command in the command console.
      - **Linux** | **UNIX** `./stopServer.sh server_name`
      - **Windows** `stopServer.bat server_name`
    - c) When prompted, enter the user ID and password of WebSphere administrative console administrator.

d) Start the application server again by entering the **startServer** command in the command console.

- **Linux** | **UNIX** `./startServer.sh server_name`
- **Windows** `startServer.bat server_name`

## Results

- The data collector is configured to monitor all instances in a profile, or, for WebSphere Application Server Liberty, a single instance or multiple instances in the same directory. To monitor more profiles or instances, repeat the configuration.
- The data collector is configured within the server instances, providing maximum monitoring.

**Known limitation:** When monitoring WebSphere Application Server Liberty, the data collector cannot generate Java Naming and Directory Interface (JNDI) events.

## What to do next

Log in to the Cloud App Management console to view monitoring data. For more information, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring or reconfiguring the data collector with full configuration utilities

Use the full configuration utilities (interactive or silent) to configure the data collector for the WebSphere Applications agent. You can also use full configuration utilities to reconfigure the data collector when it is already configured. Also, you need to use the full configuration utility to configure monitoring for WebSphere Portal Server instances.

### Before you begin

Make sure that the user access requirements are met in your environment. For instructions, see [“Checking user access requirements”](#) on page 525.

**Restriction:** To monitor WebSphere Application Server Liberty, make sure that the `<featureManager>` section is defined in the Liberty `server.xml` file. Otherwise, data collector configuration cannot add required features to load to the `server.xml` file for monitoring purpose.

Sometimes, the required features are not defined in the `<featureManager>` section of `server.xml`, but in an external `feature.xml` file. Then, the **include** element is used in `server.xml` to include the feature information from the external `xml` file. In this case, you must remove the **include** element from `server.xml` and then copy the features in the external `xml` file to the `<featureManager>` section of `server.xml`.

### About this task

The configuration and reconfiguration utilities can be found in the following directories:

- **Windows** `install_dir\dchome\7.3.0.14.09\bin`
- **Linux** | **UNIX** `install_dir/yndchome/7.3.0.14.09/bin`

### Procedure

- The configuration utility is named **config**. You might need to configure the data collector with the full configuration utility in the following cases:
  - You want to configure the data collector for the first time after the WebSphere Applications agent is installed.
  - You want to configure monitoring for WebSphere Portal Server instances.
  - You do not want to configure all application servers within the same profile at one time.
  - The data collector is not configured within the application server and you want to reconfigure it.

For information about the interactive full configuration utility, see [“Configuring the data collector interactively”](#) on page 528.

- The reconfiguration utility is named **reconfig**. You might need to reconfigure the data collector in the following cases:
  - You want to reconfigure the data collector after it is configured either interactively or silently.

For information about interactive reconfiguration utility, see [“Reconfiguring the data collector interactively”](#) on page 531.

- For silent configuration, see [“Configuring the data collector in silent mode”](#) on page 534.

### Configuring the data collector interactively

Use the interactive configuration utility (`config.sh` or `config.bat`) to configure the data collector for each application server instance that you want to monitor.

### Before you begin

If you will configure the data collector to monitor WebSphere Application Server Liberty, set the **JAVA\_HOME** system environment variable to the same JVM as the one used for the application server. For example, on a Windows system, set **JAVA\_HOME** value to `C:\Program Files\IBM\java`. Or on a Linux system, run `export JAVA_HOME=/opt/IBM/java`.

### About this task

Use the following full configuration utility to configure the data collector:

- **Windows** `install_dir\dchome\7.3.0.14.09\bin\config.bat`
- **Linux** | **UNIX** `install_dir/yndchome/7.3.0.14.09/bin/config.sh`

### Procedure

To configure the data collector by responding to prompts, complete these steps:

1. Log in to the system with the user ID that is used to install the application server.
2. Go to the `bin` directory within the `dc_home` data collector home directory.
3. Start the configuration utility by issuing the following command:

- **Windows** `config.bat`
- **Linux** | **UNIX** `./config.sh`

The configuration utility displays the IP addresses and host names of all network cards that are found on the local computer system.

4. Enter the number that corresponds to the IP address and host name. If the IP address and host name that you want to use are not on the list, enter the IP address or host name.
5. Specify the home directory of the application server that is to be monitored.
  - For traditional WebSphere Application Server, enter the number that corresponds to an auto-discovered application server home directory or specify a full path to an application server home directory.
  - For WebSphere Application Server Liberty, enter the full path to the WebSphere Application Server Liberty home directory that contains the `bin` and `servers` directories, for example `/opt/ibm/wlp`.
6. If you are configuring the data collector for WebSphere Application Server Liberty, you are prompted for the Java home directory. Specify the Java home directory that is used for the application server. For example, `/opt/IBM/java`.
7. When the configuration utility lists all profiles under the specified application server home directory, enter the number that corresponds to the application server profile that you want to configure.

- For traditional WebSphere Application Server, the configuration utility then indicates whether WebSphere Global Security is enabled for the WebSphere Application Server profile that you specified. If global security is not enabled, proceed to the Step “9” on page 529.
  - For WebSphere Application Server Liberty, proceed to Step “10” on page 529.
8. If global security is enabled for the WebSphere Application Server profile, specify whether to retrieve security settings from a client properties file. Enter 1 to allow the configuration utility to retrieve the user name and password from the appropriate client properties file. Otherwise, enter 2 to enter the user name and password.

The data collector communicates with the WebSphere Administrative Services by using the Remote Method Invocation (RMI) or the SOAP protocol. If global security is enabled for a profile, you must specify the user ID and password of a user who is authorized to log in to the WebSphere Application Server administrative console for the application server profile. Alternatively, you can encrypt the user name and password and store them in client properties files before you configure the data collector. You must use the `sas.client.props` file for an RMI connection, or the `soap.client.props` file for an SOAP connection.

9. When you are prompted for the host name of WebSphere administrative console, press Enter to accept the default or specify the host name or IP address of the WebSphere administrative console. The default value is `localhost`.

**Remember:**

- For a Network Deployment environment, enter the host name or IP address of the Deployment Manager.
  - The maximum length of the host name is 19 characters. If the value that you specify exceeds 19 characters, the host name will be truncated. For instructions about how to change the host name, see “Changing the host name used in MSN” on page 538.
10. When the configuration utility lists all the server instances that are not configured yet for data collection, select one or more application server instances from the list. Enter the number that corresponds to the application server instance to configure for data collection or enter an asterisk (\*) to configure all application server instances for data collection. To specify a subset of servers, enter the numbers that represent the servers, separated by commas. For example, `1, 2, 3`.

**Remember:**

- For a stand-alone environment, application server instances must be running during the configuration. (A WebSphere Application Server Liberty instance does not need to be running).
  - For a Network Deployment environment, the Deployment Manager must be running.
  - Ensure that the application server instances that you select are the actual servers that host the applications or services that you want to monitor.
11. In the **Integration with Agent for WebSphere Applications** section, specify that you want to integrate the data collector with the WebSphere Applications agent. You must enter 1 to select this integration option, and then press Enter.

The selected server will be registered for PMI resource monitoring.

12. If you are configuring the data collector for traditional WebSphere Application Server, specify whether you want to configure the data collector within the application server instance.
- Enter 1 to configure the data collector within the application server. With this option, the data collector is integrated with the application server, which is required for the full range of operational monitoring and diagnostics data collection. However, configuring the data collector within the application server requires restarting the application server. Also, the data collector might affect server performance.
  - Enter 2 to not to configure the data collector within the application server and proceed to Step “13” on page 530. With this option, the data collector runs as a stand-alone process and only resource monitoring can be enabled.

13. When you are prompted for the host name of the V8 monitoring agent, enter the host name or IP address of the WebSphere Applications agent or press Enter to accept the default. The default value corresponds to your choice in Step 3.

The V8 monitoring agent refers to the WebSphere Applications agent, which is installed with Cloud App Management.

14. When you are prompted for the port number of the V8 monitoring agent, enter the port number of the WebSphere Applications agent or press Enter to accept the default. The default is 63335.
15. When you are asked whether to configure V6 monitoring agent for WebSphere Applications, press Enter to accept the default for No.
16. When you are prompted for the server alias, press Enter to accept the default or enter another alias. If you are configuring several application server instances, the configuration utility prompts you for an alias for every instance.

**Important:** The alias can contain only the following characters: A-Z, a-z, underbar (\_), dash (-), and period (.). Do not use other characters in the alias.

17. When you are prompted for a port number for PMI resource monitoring, press Enter to accept the default or enter a new number. The default port is 63355.

This port is used for internal communication between components that are running on the same host. If the default is in use, you can set a different number.

18. In the **Support for transaction tracking** section, specify whether to enable transaction tracking. Enter 1 to enable support for transaction tracking. Otherwise, enter 2.

**Note:** To make sure the Transaction Tracking feature works, the JDK version must be above 1.7.

19. When you are prompted for the port number that the data collector uses to connect to the Transaction Framework Extension, press Enter to accept the default or enter another port number. The default is 5457.
20. In the **Advanced settings** section, specify whether to change the garbage collection log path. Enter 1 to select a garbage collection log path. Otherwise, enter 2 and skip to Step “22” on page 530. To use the log path that is already specified in the JVM argument of the application server, enter 2.
21. Specify the garbage collection log path. Enter a file name with its full path. For WebSphere Application Server Liberty, do not use variables in the path. The data collector automatically modifies the log file name, adding the server instance information to it.

For example, if you specify `gc.log` as the file name, the actual name is set to `profile_name.cell_name.node_name.server_name.gc.log` for every configured application server instance.

**Important:** In the garbage collection log path, you can use WebSphere variables such as `${SERVER_LOG_ROOT}`. However, do not use templates, such as `%pid`.

22. Review the summary of the data collector configuration that is to be applied to the specified application server instances. If necessary, reconfigure parts of the data collector configuration before you apply the changes.
23. Enter a to accept your changes.
24. When prompted, specify whether you want to create a backup of your current configuration. Enter 1 to create a backup of the current configuration. Otherwise, enter 2.

The configuration utility applies the changes and presents a status message to indicate that the configuration of the data collector for the profile is completed.

25. If you are configuring the data collector for traditional WebSphere Application Server, restart the application server instances or restart the agent, depending on your choice in Step “12” on page 529.
  - If you have enabled the data collector within the application server, restart the application server instances as indicated by the configuration utility.
  - If you have enabled PMI resource monitoring without enabling the data collector within the application server, restart the WebSphere Applications agent by running the following commands:

– Windows

```
cd install_dir\bin
was-agent.bat stop
was-agent.bat start
```

– Linux | UNIX

```
cd install_dir/bin
./was-agent.sh stop
./was-agent.sh start
```

The data collector configuration takes effect after the application server or agent restart.

## What to do next

- If the current user ID that is used to configure the data collector is not the same ID of the user running the application server, verify that the user ID for configuring the data collector has read and write permissions to the `runtime` and `logs` directories within the data collector home directory. These two sub-directories are created by the ID of the user running the application server when the server is restarted.
- Log in to the Cloud App Management console to view the monitoring data in the dashboards. If monitoring data are not available immediately, restart the WebSphere Applications agent by running the following commands:

– Windows

```
cd install_dir\bin
was-agent.bat stop
was-agent.bat start
```

– Linux | UNIX

```
cd install_dir/bin
./was-agent.sh stop
./was-agent.sh start
```

## Reconfiguring the data collector interactively

If you configured the data collector to monitor one or more application server instances, you can reconfigure the data collector by using the reconfiguration utility (`reconfig.sh` or `reconfig.bat`).

## Before you begin

If you will configure the data collector to monitor WebSphere Application Server Liberty, set the **JAVA\_HOME** system environment variable to the same JVM as the one used for the application server. For example, on a Windows system, set **JAVA\_HOME** value to `C:\Program Files\IBM\java`. Or on a Linux system, run `export JAVA_HOME=/opt/IBM/java`.

## About this task

Use the following full reconfiguration utility to configure the data collector:

- Windows `install_dir\dchome\7.3.0.14.09\bin\reconfig.bat`
- Linux | UNIX `install_dir/yndchome/7.3.0.14.09/bin/reconfig.sh`

**Remember:** The **reconfig** utility is not applicable in the following cases. Use the **config** configuration utility instead. Although the **config** utility warns that the server is already configured, but it still can make any required changes.

- The data collector is already configured for resource monitoring only and you want to reconfigure the data collector.
- You want to reconfigure the data collector for WebSphere Portal Server.

**Tip:** In the prompts asking for agent configuration settings, the reconfiguration utility offers the currently configured values as defaults.

## Procedure

To reconfigure the data collector by responding to prompts, complete these steps:

1. Log in to the system with the user ID that is used to install the application server.
2. Go to the `bin` directory within the `dc_home` data collector home directory.
3. Start the reconfiguration utility by issuing the following command:

- **Windows** `reconfig.bat`
- **Linux** | **UNIX** `./reconfig.sh`

**Tip:** Running this reconfiguration utility has the same effect as running the `config.bat` script with the `-reconfig` argument on Windows systems or the `config.sh` script with the `-reconfig` argument on Linux or AIX systems.

The reconfiguration utility displays the IP addresses of all network cards that are found on the local computer system.

4. Enter the number that corresponds to the IP address to use.  
The reconfiguration utility displays all application server instances for which the data collector is configured on this host, and prompts you to select one or more application server instances from the list.
5. Select one or more application server instances from the list. Enter the number that corresponds to the application server instance to reconfigure for data collection or enter an asterisk (\*) to reconfigure all application server instances for data collection. To specify a subset of servers, enter the numbers, separated by commas, that represent the servers. For example: `1, 2, 3`.

### Remember:

- For a stand-alone environment, application server instances must be running during the configuration. (A WebSphere Application Server Liberty instance does not need to be running).
  - For a Network Deployment environment, the Deployment Manager must be running.
  - Ensure that the application server instances that you select are the actual servers that host the applications or services that you want to monitor.
6. In the **Integration with Agent for WebSphere Applications** section, specify that you want to integrate the data collector with the WebSphere Applications agent. You must enter `1` to select this integration option, and then press Enter.
  7. If you are configuring the data collector for traditional WebSphere Application Server, specify whether you want to configure the data collector within the application server instance.
    - Enter `1` to configure the data collector within the application server. With this option, the data collector is integrated with the application server, which is required for the full range of operational monitoring and diagnostics data collection. However, configuring the data collector within the application server requires restarting the application server. Also, the data collector might affect server performance.
    - Enter `2` to not to configure the data collector within the application server and process to Step “8” on page 532. With this option, the data collector runs as a stand-alone process and only PMI resource monitoring can be enabled.
  8. When you are prompted for the host name, enter the host name or IP address of the WebSphere Applications agent or press Enter to accept the default. The default value corresponds to your choice in Step “4” on page 532.
  9. When you are prompted for the port number, enter the port number of the monitoring agent or press Enter to accept the default. The default is `63335`.
  10. When you are asked whether to configure V6 monitoring agent for WebSphere Applications, press Enter to accept the default for No.



11. When you are prompted for the server alias, press Enter to accept the default or enter another alias. If you are configuring several application server instances, the configuration utility prompts you for an alias for every instance.

**Important:** The alias can contain only the following characters: A-Z, a-z, underbar (\_), dash (-), and period (.). Do not use other characters in the alias.

12. When you are prompted for a port number for PMI resource monitoring, press Enter to accept the default or enter a new number. The default port is 63355.

This port is used for internal communication between components that are running on the same host. If the default is in use, you can set a different number.

13. Specify whether to integrate the data collector with Application Performance Diagnostics Lite. Press Enter to accept the default for No.

14. In the **Advanced settings** section, specify whether to change the garbage collection log path.

Enter 1 to select a garbage collection log path. Otherwise, enter 2 and skip to Step “16” on page 533. To use the log path that is already specified in the JVM argument of the application server, enter 2.

15. Specify the garbage collection log path. Enter a file name with its full path. For WebSphere Application Server Liberty, do not use variables in the path. The data collector automatically modifies the log file name, adding the server instance information to it.

For example, if you specify `gc.log` as the file name, the actual name is set to `profile_name.cell_name.node_name.server_name.gc.log` for every configured application server instance.

**Important:** In the garbage collection log path, you can use WebSphere variables such as `{SERVER_LOG_ROOT}`. However, do not use templates, such as `%pid`.

16. Review the summary of the data collector configuration that is to be applied to the specified application server instances. Reconfigure parts of the data collector configuration before you apply the changes, if required.

17. Enter a to accept your changes.

18. When prompted, specify whether you want to create a backup of your current configuration. Enter 1 to create a backup of the current configuration. Otherwise, enter 2.

The configuration utility applies the changes and presents a status message to indicate that the configuration of the data collector for the profile is completed.

19. If you are configuring the data collector for traditional WebSphere Application Server, restart the application server instances or restart the agent, depending on your choice in Step “7” on page 532.

- If you have enabled the data collector within the application server, restart the application server instances as indicated by the configuration utility.
- If you have enabled PMI resource monitoring without enabling the data collector within the application server, restart the WebSphere Applications agent by running the following commands:

– Windows

```
cd install_dir\bin
was-agent.bat stop
was-agent.bat start
```

– Linux | UNIX

```
cd install_dir/bin
./was-agent.sh stop
./was-agent.sh start
```

The data collector configuration takes effect after the application server or agent restart.

## Configuring the data collector in silent mode

If you want to configure many application server instances, it might be more convenient to configure the data collector in silent mode.

### About this task

When you configure the data collector in silent mode, you first specify configuration options in a properties file. A sample properties file, `sample_silent_config.txt`, is packaged with the configuration utility. The file is available in the following directories:

- **Linux** | **UNIX** `install_dir/yndchome/7.3.0.14.09/bin`
- **Windows** `install_dir\dchome\7.3.0.14.09\bin`

For detailed information about each available configuration property in this file, see [“Properties file for silent configuration of data collector”](#) on page 535.

### Procedure

Complete the following steps to perform a silent configuration:

1. Specify configuration options in the properties file. You can copy the sample properties file and change the required options.
2. Set the location of the Java home directory before you run the utility.  
For example:

- **Windows**

```
set JAVA_HOME=C:\Program Files\IBM\WebSphere\AppServer80\java
```

- **Linux** | **UNIX**

```
export JAVA_HOME=/opt/IBM/AppServer80/java
```

**Important:** If you are configuring monitoring for WebSphere Application Server Liberty, you must use same JVM version as the one used for the application server. Otherwise, the monitoring might fail.

3. Go to the following directory:

- **Linux** | **UNIX** `install_dir/yndchome/7.3.0.14.09/bin`
- **Windows** `install_dir\dchome\7.3.0.14.09\bin`

4. Run the configuration command to configure the data collector in silent mode.

**Tip:** If the `wsadmin` user was used to install the application server, run the `config` utility either as the `wsadmin` user or with root user privileges.

- **Windows** Run the following command as the administrator who installed the WebSphere Application Server.

```
config.bat -silent path_to_silent_file
```

- **Linux** | **UNIX** Run the following command with root user privileges.

```
config.sh -silent path_to_silent_file
```

where, *full\_path\_to\_silent\_file* is the path to the silent `.txt` file.

5. After configuring the data collector to monitor application server instances, if you have enabled the data collector within the application server, you must restart the instances. The data collector configuration takes effect when the application server instances are restarted.
6. If you have enabled PMI resource monitoring without enabling the data collector within the application server, you might need to restart the WebSphere Applications agent to start the monitoring. If

monitoring data is not available immediately, restart the monitoring agent by running the following commands:

- **Windows**

```
cd install_dir\bin
was-agent.bat stop
was-agent.bat start
```

- **Linux | UNIX**

```
cd install_dir/bin
./was-agent.sh stop
./was-agent.sh start
```

### What to do next

After silent configuration, to reconfigure the data collector, you have two options:

- Reconfigure it interactively by using the **reconfig** reconfiguration utility. For instructions, see [“Reconfiguring the data collector interactively”](#) on page 531.
- Unconfigure it silently and then use the same procedure to configure it silently again. For instructions, see [“Unconfiguring the data collector in silent mode”](#) on page 217.

### Properties file for silent configuration of data collector

To silently configure the data collector, you first specify configuration options in a properties file and then run the configuration utility.

When you create your properties file, keep in mind the following considerations:

- A line in the file that starts with a number sign (#) is treated as a comment, and is not processed. If the number sign is used elsewhere in the line, it is not considered to be the start of a comment.
- Each property is described on a separate line, in the following format: *property = value*.

#### **property**

Name of property. The list of valid properties that you can configure is shown in [Table 75 on page 535](#).

#### **value**

Value of the property. Default values for some properties are already provided. You can delete default values to leave property values blank or empty. An empty value is treated as if the property is not specified, as opposed to using the default value. If you want to use default values, you can comment out the property in the file.

- Passwords are in plain text.
- Properties and their values are case-sensitive.

[Table 75 on page 535](#) describes the properties that are available when configuring the data collector in silent mode.

**Important:** If you are configuring the data collector for a WebSphere Application Server Liberty instance, some of the properties are not used.

| <i>Table 75. Available properties for running the configuration utility in silent mode</i> |                                                                                                          |
|--------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>Property</b>                                                                            | <b>Comment</b>                                                                                           |
| default.hostip                                                                             | If the computer system uses multiple IP addresses, specify the IP address for the data collector to use. |
| <b>Support for transaction tracking</b>                                                    |                                                                                                          |

Table 75. Available properties for running the configuration utility in silent mode (continued)

| Property                                                                                                                                                                                                                                                       | Comment                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ttapi.enable                                                                                                                                                                                                                                                   | Specifies whether the data collector supports transaction tracking. Valid values are True and False.<br><b>Remember:</b> You must set it to False because transaction tracking is not supported by Cloud App Management.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| ttapi.host                                                                                                                                                                                                                                                     | Specifies the host of the Transaction Framework Extension, which is the component of the WebSphere Applications agent that gathers metrics from the data collector. Use the local host value, 127.0.0.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| ttapi.port                                                                                                                                                                                                                                                     | Specifies the port of the Transaction Framework Extension. Use 5457.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>PMI resource and data collector monitoring</b>                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| The selected server is always configured for resource (PMI) monitoring, without any changes to the application server. This monitoring option provides limited metrics, but does not require restarting the application server and can not affect performance. |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| tema.appserver                                                                                                                                                                                                                                                 | Specifies whether you want to configure the data collector within the application server instance. The data collector within the application server instance is required for the full range of metrics in WebSphere Applications agent and for integration with any other products. However, configuring the data collector requires restarting the application server. Also, the data collector might affect server performance. Valid values are True and False.<br><br>When this parameter is set to False, diagnostics and transaction tracking features are not available, and only resource monitoring data is collected.                                                                                                                                                                                                            |
| tema.jmxport                                                                                                                                                                                                                                                   | TCP/IP port number for resource monitoring. The port is used for internal communication between components running on the same host. The default port is 63355; if this port is in use, you can set a different number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Integration of the data collector with the monitoring agent component of WebSphere Applications agent</b>                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| temaconnect                                                                                                                                                                                                                                                    | Specifies whether the data collector connects to the monitoring agent component of WebSphere Applications agent. Valid values are True and False.<br><b>Important:</b> You must use the True value to use the WebSphere Applications agent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| tema.appserver                                                                                                                                                                                                                                                 | Specifies whether you want to configure the data collector within the application server instance. The data collector within the application server instance is required for the full range of metrics in the WebSphere Applications agent and for integration with any other products. However, it requires restarting the application server. Also, the data collector might affect server performance. Valid values are True and False.<br><br>If this parameter is set to False, the configuration parameters for integrating data collector with products other than WebSphere Applications agent are disregarded, as well as the following tema.host and tema.port parameters. When this parameter is set to False, diagnostics and transaction tracking features are not available, and only resource monitoring data is collected. |

Table 75. Available properties for running the configuration utility in silent mode (continued)

| Property                                                     | Comment                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| tema.host                                                    | Specifies the fully qualified host name or IP address of the monitoring agent component of WebSphere Applications agent. Use the local host address (127.0.0.1).                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| tema.port                                                    | Specifies the port number of the monitoring agent component of WebSphere Applications agent. Do not change the default value of 63335.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| tema.jmxport                                                 | TCP/IP port number for resource monitoring. The port is used for internal communication between components running on the same host. The default port is 63355; if this port is in use, you can set a different number.                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>WebSphere Application Server backup</b>                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| was.backup.configuration                                     | Specifies whether to back up the current configuration of the WebSphere Application Server configuration before applying the new configuration. Valid values are True and False.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| was.backup.configuration.dir                                 | Specifies the location of the backup directory.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Advanced configuration settings</b>                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| was.gc.custom.path                                           | Specifies whether to set a custom path for the Garbage Collection log.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| was.gc.file                                                  | <p>Specifies the path to the custom Garbage Collection log. Set this value to a file name with its full path. The data collector automatically modifies the log file name, adding the server instance information to it. For example, if you specify gc.log as the file name, the actual name is set to <i>profile_name.cell_name.node_name.server_name.gc.log</i> for every configured application server instance.</p> <p><b>Important:</b> In the Garbage Collection log path, you can use WebSphere variables, such as <code>\${SERVER_LOG_ROOT}</code>. However, do not use templates, such as <code>%pid</code>.</p> |
| <b>WebSphere Administrative Services connection settings</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| was.wsadmin.connection.host                                  | <p>Specifies the name of the host to which the wsadmin tool is connecting. In a Network Deployment environment, specify the wsadmin connection to the Deployment Manger. In a stand-alone environment, specify the wsadmin connection to the server.</p> <p><b>Remember:</b> If the WebSphere Administrative console is on the same system, the value of localhost will be used for connection. However, in some cases, localhost is not allowed for communication due to system network or security settings. In that case, you must specify this parameter in the silent response file.</p>                              |
| was.wsadmin.connection.type                                  | Specifies the port that the wsadmin tool must use to connect to the WebSphere Application Server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| was.wsadmin.connection.port                                  | Specifies the port that the wsadmin tool must use to connect to the WebSphere Application Server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>WebSphere Application Server global security settings</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| was.wsadmin.username                                         | Specifies the user ID of a user who is authorized to log in to the WebSphere Application Server administrative console. This user must have the agent role on the application server.                                                                                                                                                                                                                                                                                                                                                                                                                                      |

| <i>Table 75. Available properties for running the configuration utility in silent mode (continued)</i> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Property</b>                                                                                        | <b>Comment</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| was.wsadmin.password                                                                                   | Specifies the password that corresponds to the user specified in the was.wsadmin.username property.                                                                                                                                                                                                                                                                                                                                                                |
| was.client.props                                                                                       | Specifies whether to retrieve security settings from a client properties file. Possible values are True and False.                                                                                                                                                                                                                                                                                                                                                 |
| <b>WebSphere Application Server settings</b>                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| was.appserver.profile.name                                                                             | Specifies the name of the application server profile that you want to configure. Not used for WebSphere Application Server Liberty.                                                                                                                                                                                                                                                                                                                                |
| was.appserver.home                                                                                     | Specifies the WebSphere Application Server home directory.                                                                                                                                                                                                                                                                                                                                                                                                         |
| was.appserver.cell.name                                                                                | Specifies the WebSphere Application Server cell name. Not used for WebSphere Application Server Liberty.                                                                                                                                                                                                                                                                                                                                                           |
| was.appserver.node.name                                                                                | Specifies the WebSphere Application Server node name. Not used for WebSphere Application Server Liberty.                                                                                                                                                                                                                                                                                                                                                           |
| <b>WebSphere Application Server runtime instance settings</b>                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| was.appserver.server.name                                                                              | Specifies the application server instance within the application server profile to configure.<br><b>Tip:</b> <ul style="list-style-type: none"> <li>• The silent response file can have multiple instances of this property</li> <li>• When adding a second server, uncomment the second server (this is, #[SERVER]) and add the server name.</li> </ul>                                                                                                           |
| tema.serveralias                                                                                       | Specifies the name of the node in monitoring user interface that contains the monitoring information for this application server instance. The default is the node name combined with the server name.<br><b>Important:</b> The alias can contain only the following characters: A-Z, a-z, underbar (_), dash (-), and period (.). Do not use other characters in the alias.<br><b>Tip:</b> The silent response file can have multiple instances of this property. |

### Changing the host name used in MSN

The managed system name (MSN) is the instance name that you see on the Resources Dashboard. The host name used in MSN has a maximum limit of 19 characters. If the length exceeds 19 characters, the instance property value is truncated in the Resources Dashboard. Then you must change the host name.

### About this task

When the WebSphere Applications agent is started, it registers the following MSN for each agent instance:

*serveralias:hostname:KYNS*

Where:

- *serveralias* is the alias that you assign to the application server during data collector configuration.
- *hostname* is the name of the host where the agent is running.
- KYNS is the fixed string that identifies the WebSphere Applications agent.

### Important:

- MSN has a maximum length limit of 32 characters.
- KYNS is fixed and cannot be changed.

- The maximum length of *hostname* is 19 characters. If the length exceeds 19 characters, the instance property value is truncated in the Resources Dashboard.
- The maximum length of *serveralias* equals 26 minus the length of *hostname*. If the length exceeds, the instance property value is truncated in the Resources Dashboard.
- Any truncation of the MSN attributes causes the incorrect display of resources names and property values.

If the length of *hostname* or *serveralias* exceeds, the specified string is truncated.

To avoid the truncation issue, you can follow the steps to change the *hostname* value.

## Procedure

1. Stop the WebSphere Applications agent.
2. Open the following agent environment file with a text editor. If the file does not exist, create it by yourself.
  - Linux: *agent\_install\_dir/config/yn.environment*
  - Windows: *agent\_install\_dir/Config/KYNNENV*
3. Add the following variable to set the new host name in the agent environment file, where *newhostname* is the new host name with a maximum length of 19 characters.

```
CTIRA_HOSTNAME=newhostname
```

4. Back up the following xml file in the *agent\_install\_dir/config/* directory and then remove the original one:

```
hostname_yn_wasversion.cellname.nodename.profilename.servername.xml
```

For example,

```
/opt/ibm/apm/agent/config/tivvm123_yn_was85.tivvm123Node01Cell.tivvm123Node01.AppSrv01.server1.xml
```

### Remember:

If you configured multiple data collector instances for the WebSphere Applications agent, there will be multiple xml files. Each application server has its own xml file. You must remove all the xml files.

5. Open the *agent\_install\_dir/config/hostname\_yn.xml* file and make the following changes:
  - In the `<!DOCTYPE AgentConfig[ ]>` section, remove the following entry:

```
<!ENTITY wasversion.cellname.nodename.profilename.servername SYSTEM "hostname_yn_was_version.cellname.nodename.profilename.servername.xml">
```

For example, remove:

```
<!ENTITY was85.tivvm123Node01Cell.tivvm123Node01.AppSrv01.server1 SYSTEM "tivvm123_yn_was85.tivvm123Node01Cell.tivvm123Node01.AppSrv01.server1.xml">
```

- Between the `</defaultServerSettings>` and `</AgentConfig>` tags, remove the `"&was_version.cellname.nodename.profilename.servername"` line.

For example, remove `&was85.tivvm123Node01Cell.tivvm123Node01.AppSrv01.server1;`

6. Start the WebSphere Applications agent. An error might occur, saying that data is not sent due to the missing xml file. You can ignore this error and try starting the agent again.
7. Stop the WebSphere Applications agent. The previously removed xml file in Step 4 is created again.
8. Go to the *agent\_install\_dir/logs* directory and verify that the new host name is applied to the data collector instance in the log files.

For example, on a Linux system, you can issue `grep <newhostname>:KYNS' *asf*` to check whether there is any result returned.

9. Start the WebSphere Applications agent again. The agent instance will work under the new MSN.

### Example

A real example of modifying the `hostname_yn.xml` file, for example, `tivvm123_yn.xml` in Step 5.

Original content:

```
<!DOCTYPE AgentConfig [
  <!ENTITY was85.tivvm123Node01Cell.tivvm123Node01.AppSrv01.server1 SYSTEM
"tivvm123_yn_was85.tivvm123Node01Cell.tivvm123Node01.AppSrv01.server1.xml">
]>
<AgentConfig version="07.30.14.000">
  <altNodeId>Primary</altNodeId>
  <port>63335</port>
  <host>127.0.0.1</host>
  <maxAgentLogMsgs>100</maxAgentLogMsgs>
  <migrationData>
  </migrationData>
  <wXSConfig>
    <aliases-catalog>
      <dictionary name="zones">
      </dictionary>
      <dictionary name="grids">
      </dictionary>
      <dictionary name="mapSets">
      </dictionary>
      <dictionary name="maps">
      </dictionary>
      <dictionary name="hosts">
      </dictionary>
      <dictionary name="catalogs">
      </dictionary>
      <dictionary name="containers">
      </dictionary>
      <dictionary name="coreGroups">
      </dictionary>
      <dictionary name="domains">
      </dictionary>
    </aliases-catalog>
  </wXSConfig>
  <xDAgentConfig>
    <reconnectDelaySeconds>60</reconnectDelaySeconds>
  </xDAgentConfig>
  <defaultServerSettings use-agent-settings="true">
    <resourceMonitoringMethod>ON_DEMAND</resourceMonitoringMethod>
    <resourceMonitoringLevel>ALL</resourceMonitoringLevel>
    <resourceMonitoringEnabled>true</resourceMonitoringEnabled>
    <requestMonitoringLevel>L1</requestMonitoringLevel>
    <gCMaxLogEvents>100</gCMaxLogEvents>
    <gCMonitoringEnabled>true</gCMonitoringEnabled>
    <requestMonitoringMethod>FIXED_INTERVAL</requestMonitoringMethod>
    <dataCollectorSettings>
    </dataCollectorSettings>
    <advancedSettings>
      <appSrvLogScanEvtEmitThreshold>1</appSrvLogScanEvtEmitThreshold>
      <resourceFixIntervalTime>60</resourceFixIntervalTime>
      <reqSampleRate>2</reqSampleRate>
      <requestOnDemandSampleAge>30</requestOnDemandSampleAge>
      <hangThreadDetectionTimeout>300</hangThreadDetectionTimeout>
      <resourceOnDemandSampleAge>30</resourceOnDemandSampleAge>
      <hungThreadsMonitoringEnabled>true</hungThreadsMonitoringEnabled>
      <requestFixIntervalTime>60</requestFixIntervalTime>
      <appSrvLogMaxMsgs>100</appSrvLogMaxMsgs>
      <appSrvLogScanIntervalTime>300</appSrvLogScanIntervalTime>
      <gCLogScanIntervalTime>60</gCLogScanIntervalTime>
    </advancedSettings>
    <applicationMonitoringSettings>
      <respTimeAutoTresholdGoodZoneProjection>150
      </respTimeAutoTresholdGoodZoneProjection>
      <resUsageFairThreshold>40</resUsageFairThreshold>
      <compRateFair>99</compRateFair>
      <respTimeAutoTresholdDeviation>200</respTimeAutoTresholdDeviation>
      <respTimeAutoTresholdSelection>50</respTimeAutoTresholdSelection>
      <resUsageMonitoringThreshold>25</resUsageMonitoringThreshold>
      <respTimeAutoTresholdFairZoneProjection>300
```



```

        </respTimeAutoTresholdFairZoneProjection>
        <resUsageBadThreshold>80</resUsageBadThreshold>
        <requestMonitoringMode>APPLICATION</requestMonitoringMode>
        <complRateBad>95</complRateBad>
    </applicationMonitoringSettings>
</defaultServerSettings>
&was85.tivvm123Node01Cell.tivvm123Node01.AppSrv01.server1;
</AgentConfig>

```

Modified content:

```

<!DOCTYPE AgentConfig [
]>
<AgentConfig version="07.30.14.000">
  <altNodeId>Primary</altNodeId>
  <port>63335</port>
  <host>127.0.0.1</host>
  <maxAgentLogMsgs>100</maxAgentLogMsgs>
  <migrationData>
  </migrationData>
  <wXSConfig>
    <aliases-catalog>
      <dictionary name="zones">
      </dictionary>
      <dictionary name="grids">
      </dictionary>
      <dictionary name="mapSets">
      </dictionary>
      <dictionary name="maps">
      </dictionary>
      <dictionary name="hosts">
      </dictionary>
      <dictionary name="catalogs">
      </dictionary>
      <dictionary name="containers">
      </dictionary>
      <dictionary name="coreGroups">
      </dictionary>
      <dictionary name="domains">
      </dictionary>
    </aliases-catalog>
  </wXSConfig>
  <xDAgentConfig>
    <reconnectDelaySeconds>60</reconnectDelaySeconds>
  </xDAgentConfig>
  <defaultServerSettings use-agent-settings="true">
    <resourceMonitoringMethod>ON_DEMAND</resourceMonitoringMethod>
    <resourceMonitoringLevel>ALL</resourceMonitoringLevel>
    <resourceMonitoringEnabled>true</resourceMonitoringEnabled>
    <requestMonitoringLevel>L1</requestMonitoringLevel>
    <gCMaxLogEvents>100</gCMaxLogEvents>
    <gCMonitoringEnabled>true</gCMonitoringEnabled>
    <requestMonitoringMethod>FIXED_INTERVAL</requestMonitoringMethod>
    <dataCollectorSettings>
    </dataCollectorSettings>
    <advancedSettings>
      <appSrvLogScanEvtEmitThreshold>1</appSrvLogScanEvtEmitThreshold>
      <resourceFixIntervalTime>60</resourceFixIntervalTime>
      <reqSampleRate>2</reqSampleRate>
      <requestOnDemandSampleAge>30</requestOnDemandSampleAge>
      <hangThreadDetectionTimeout>300</hangThreadDetectionTimeout>
      <resourceOnDemandSampleAge>30</resourceOnDemandSampleAge>
      <hungThreadsMonitoringEnabled>true</hungThreadsMonitoringEnabled>
      <requestFixIntervalTime>60</requestFixIntervalTime>
      <appSrvLogMaxMsgs>100</appSrvLogMaxMsgs>
      <appSrvLogScanIntervalTime>300</appSrvLogScanIntervalTime>
      <gCLogScanIntervalTime>60</gCLogScanIntervalTime>
    </advancedSettings>
  </applicationMonitoringSettings>
    <respTimeAutoTresholdGoodZoneProjection>150
    </respTimeAutoTresholdGoodZoneProjection>
    <resUsageFairThreshold>40</resUsageFairThreshold>
    <compRateFair>99</compRateFair>
    <respTimeAutoTresholdDeviation>200</respTimeAutoTresholdDeviation>
    <respTimeAutoTresholdSelection>50</respTimeAutoTresholdSelection>
    <resUsageMonitoringThreshold>25</resUsageMonitoringThreshold>
    <respTimeAutoTresholdFairZoneProjection>300
    </respTimeAutoTresholdFairZoneProjection>
    <resUsageBadThreshold>80</resUsageBadThreshold>
    <requestMonitoringMode>APPLICATION</requestMonitoringMode>

```

```
<complRateBad>95</complRateBad>
</applicationMonitoringSettings>
</defaultServerSettings>
</AgentConfig>
```

### Changing OpenTracing sampling configuration

If you enable Transaction Tracking when configuring the data collector for WebSphere Applications agent, you can change the OpenTracing sampling configuration.

#### Before you begin

Ensure that you enable Transaction Tracking for the WebSphere Applications agent. For more information, see [Configuring or reconfiguring the data collector with full configuration utilities](#).

#### About this task

You can change the OpenTracing sampling configuration by editing the following parameters in the property file.

- JAEGER\_SAMPLER\_TYPE
- JAEGER\_SAMPLER\_PARAM

The default sampler type is probabilistic, and the default sampler param is 0.01, which means that 1 in 100 traces will be sampled. You can set it to other values. For more information, see [Sampling](#).

#### Procedure

1. Navigate to the folder where the `ttdc.properties` file is located, for example, `dc_home/runtime/appserver_version.node_name.server_name/ttdc.properties`. If the folder does not exist, use `dc_home/ttdc/etc/ttdc.properties`.
2. Open the `ttdc.properties` file and update the values for `JAEGER_SAMPLER_TYPE` and `JAEGER_SAMPLER_PARAM`.

Example:

```
JAEGER_SAMPLER_TYPE=probabilistic
JAEGER_SAMPLER_PARAM=0.1
```

3. Restart your application server.

## Manually configure the data collector if the configuration utilities fail

If you cannot use the provided configuration utility to configure the data collector for WebSphere Applications agent, you can manually configure the data collector in the WebSphere Administrative Console.

#### Before you begin

- Install the WebSphere Applications agent.
- Get to know the data collector home directory, which is required by the data collector configuration. The default is `/opt/ibm/apm/agent/yndchome/7.3.0.14.09` on Linux and AIX systems or `C:\IBM\APM\dchome\7.3.0.14.09` on Windows systems.
- If you want to configure the data collector for a Liberty server, get to know the Liberty server home directory. For example, `/opt/ibm/was/liberty/usr/servers/defaultServer`.
- Make sure that a file named `itcam_wsBundleMetaData.xml` exists in the `dc_home/runtime/wsBundleMetaData` folder and it contains the following content. If the folder or the file does not exist, manually create it.

**Remember:** The `plugins_dir_within_dc_home` value must be set to the absolute path of the `plugins` folder within the data collector home directory. The default is `/opt/ibm/apm/agent/yndchome/`

7.3.0.14.09/plugins on Linux and AIX systems or C:\IBM\APM\dchome  
\7.3.0.14.09/plugins on Windows systems.

```
<bundles>
  <directory path="plugins_dir_within_dc_home">
    <bundle>com.ibm.tivoli.itcam.bundlemanager_7.2.0.jar</bundle>
  </directory>
  <directory path="plugins_dir_within_dc_home">
    <bundle>com.ibm.tivoli.itcam.classicsca_7.2.0.jar</bundle>
  </directory>
  <directory path="plugins_dir_within_dc_home">
    <bundle>com.ibm.tivoli.itcam.toolkitsca.classicsca_7.2.0.jar</bundle>
  </directory>
</bundles>
```

## About this task

### Important:

- You must make manual changes to the WebSphere Application Server configuration for data collectors as the WebSphere administrative user.
- You must be an experienced WebSphere administrator to make manual changes to the WebSphere Application Server for data collection. Any error in the manual configuration change can result in the application server not starting.
- After you manually configure the data collector to monitor application server instances, you cannot use the unconfiguration utility to unconfigure the data collector. You must manually unconfigure the data collector instead.

### Procedure

- To manually configure the data collector for the WebSphere application server, see [“Manually configuring data collector for WebSphere Application Server traditional”](#) on page 543.
- To manually configure the data collector for the Liberty server, see [“Manually configuring the data collector for WebSphere Application Server Liberty”](#) on page 545.

## Manually configuring data collector for WebSphere Application Server traditional

### Procedure

1. Log in to the WebSphere Administrative Console as the administrator.
2. In the navigation pane, click **Servers**, expand **Server Types** and click **WebSphere application servers**.
3. Under the **Server Infrastructure** section in the Configuration tab, expand **Java and Process Management** and click **Process Definition**.
4. Under the **Additional Properties** section, click **Java Virtual Machine**.
5. In the **Generic JVM arguments** field, add the following entries.

```
-agentlib:am_ibm_16=${WAS_SERVER_NAME} -Xbootclasspath/p:${ITCAMDCHOME}/  
toolkit/lib/bcm-bootstrap.jar -Djava.security.policy=${ITCAMDCHOME}/itcamdc/  
etc/datacollector.policy -verbosegc
```

When you add the entries, take note of the following:

- All entries must be on a single line.
  - Separate different arguments by spaces before the minus sign (-), and do not use spaces anywhere else.
6. Click **Apply** and then save the changes to the master configuration.
    - If you are not under a Network Deployment environment, click **Save**.
    - If you are under a Network Deployment environment, ensure that **Synchronize changes with Nodes** is selected in the **Console preferences** options and then click **Save**.

7. In the navigation pane, click **Servers**, expand **Server Types**, click **WebSphere application servers** and then click the server name.
8. In the Configuration tab, go to **Server Infrastructure > Java and Process Management > Process Definition > Environment Entries**.
9. Depending on the operating system, the hardware platform, and the application server JVM, set the following environment entry.

*Table 76. Environment entry*

Platform	Environment entry name	Environment entry value
AIX R6.1 (64-bit JVM)	LIBPATH	/lib:\${ITCAMDCHOME}/ toolkit/lib/aix536
AIX R7.1 (64 bit JVM)	LIBPATH	/lib:\${ITCAMDCHOME}/ toolkit/lib/aix536
Linux Intel R2.6 (32-bit JVM)	LD_LIBRARY_PATH	/lib:\${ITCAMDCHOME}/ toolkit/lib/li6263
Linux x86_64 R2.6 (64-bit JVM)	LD_LIBRARY_PATH	/lib:\${ITCAMDCHOME}/ toolkit/lib/lx8266
Windows (32-bit JVM)	PATH	/lib;\${ITCAMDCHOME}/ toolkit/lib/win32
Windows (64-bit JVM)	PATH	/lib;\${ITCAMDCHOME}/ toolkit/lib/win64

10. Click **Apply** and then save the changes to the master configuration.
  - If you are not under a Network Deployment environment, click **Save**.
  - If you are under a Network Deployment environment, ensure that **Synchronize changes with Nodes** is selected in the **Console preferences** options and then click **Save**.
11. In the navigation pane, click **Environment > WebSphere Variables**.
12. Specify the scope to appropriate server level and add the *ITCAMDCHOME* variable. Set the *ITCAMDCHOME* variable value to the data collector home directory. For example, `/opt/ibm/apm/agent/yndchome/7.3.0.14.09`.
13. Click **Apply** and then save the changes to the master configuration.
  - If you are not under a Network Deployment environment, click **Save**.
  - If you are under a Network Deployment environment, ensure that **Synchronize changes with Nodes** is selected in the **Console preferences** options and then click **Save**.
14. Open the following file with a text editor:
  - **Linux** | **UNIX**  
`agent_install_dir/dchome/7.3.0.14.09/runtime/DCManualInput.txt`
  - **Windows**  
`agent_install_dir\yndchome\7.3.0.14.09\runtime\DCManualInput.txt`
15. Change the `am.camtoolkit.gpe.dc.operation.mode` line as follows and save your changes.
 

```
am.camtoolkit.gpe.dc.operation.mode=WR
```
16. Restart the application server.

## What to do next

After you manually configure the data collector, you cannot use the provided `unconfig` utility to unconfigure the data collector. Manually unconfigure the data collector instead. For instructions, see [“Manually unconfigure the data collector”](#) on page 219.

## Manually configuring the data collector for WebSphere Application Server Liberty

### Procedure

1. Navigate to the Liberty server home directory. For example, `/opt/ibm/wlp/usr/servers/defaultServer`.
2. Edit the `jvm.options` file by adding the following parameters, where `dc_home` is the data collector home directory and `server_name` is the Liberty server name.. If the `jvm.options` file does not exist, create it with a text editor.

```
-agentlib:am_ibm_16=server_name
-Xbootclasspath/p:dc_home/toolkit/lib/bcm-bootstrap.jar
-Djava.security.policy=dc_home/itcamdc/etc/datacollector.policy
-verbosegc
```

When you add the entries, take note of the following things:

- Each entry must be on a single line.
  - Replace `server_name` with the actual Liberty server name. For example, `defaultServer`.
  - Replace `dc_home` with the actual data collector home directory. For example, `/opt/ibm/apm/agent/yndchome/7.3.0.14.09`.
3. Open the `server.env` file in the same directory and add the following path to the environment entry according to the operating system, where `dc_home` is the data collector home directory. If the `server.env` file does not exist, create it with a text editor.

Platform	Environment entry name	Environment entry value
AIX R6.1 (64-bit JVM)	LIBPATH	<code>/lib:<i>dc_home</i>/toolkit/lib/aix536</code>
AIX R7.1 (64 bit JVM)	LIBPATH	<code>/lib:<i>dc_home</i>/toolkit/lib/aix536</code>
Linux x86_64 R2.6 (64-bit JVM)	LD_LIBRARY_PATH	<code>/lib:<i>dc_home</i>/toolkit/lib/ix86266</code>
Linux Intel R2.6 (32-bit JVM)	LD_LIBRARY_PATH	<code>/lib:<i>dc_home</i>/toolkit/lib/li6263</code>
Windows (32-bit JVM)	PATH	<code>/lib;<i>dc_home</i>/toolkit/lib/win32</code>
Windows (64-bit JVM)	PATH	<code>/lib;<i>dc_home</i>/toolkit/lib/win64</code>

4. Open the `server.xml` file in the same directory and add the following lines to enable the monitoring feature:

```
<featureManager>
  <feature>webProfile-7.0</feature>
  <feature>monitor-1.0</feature>
  <feature>usr:itcam-730.147</feature>
</featureManager>
```

5. Open the following file with a text editor:

- **Linux** | **UNIX**

```
agent_install_dir/dchome/7.3.0.14.09/runtime/DCManualInput.txt
```

- **Windows**

```
agent_install_dir\yndchome\7.3.0.14.09\runtime\DCManualInput.txt
```

6. Change the `am.camtoolkit.gpe.dc.operation.mode` line as follows and save your changes.

```
am.camtoolkit.gpe.dc.operation.mode=WR
```

7. Restart the Liberty server.

### What to do next

After you manually configure the data collector, you cannot use the provided `unconfig` utility to unconfigure the data collector. Manually unconfigure the data collector instead. For instructions, see [“Manually unconfigure the data collector”](#) on page 219.

## Restoring the application server configuration from a backup

If you configured a stand-alone application server instance for data collection either manually or with the configuration or migration utility and the application server fails to start, you must restore the application server configuration from a backup. If you did not create a backup, contact IBM Support.

### About this task

In a Network Deployment environment, if you configured an application server instance for data collection manually or with the configuration or migration utility and the application server fails to start, you have the following options:

- You can restore the application server configuration from a backup configuration. If you did not create a backup, contact IBM Support.
- You can manually unconfigure the data collector. The Deployment Manager and the Node Agent on the application server must be running. For more information, see [“Manually removing data collector configuration from an application server instance”](#) on page 221.

This section applies only to the Windows, UNIX, and Linux operating systems.

### Procedure

To apply the backup configuration by using the **restoreConfig** command, use one of the following procedures:

- In a non-Network Deployment environment, complete the following steps:

a) Locate your backup configuration file.

The default directory is `dc_home/data`. If several backup files are present, check the modification date and time of the file. It must be the date and time of the failed configuration. If you did not complete any other data collector configurations on the same host after the failed one, use the most recent file in the directory.

b) Stop all instances of the application server.

c) Run the **restoreConfig** command from the `appserver_home/profiles/profile_name/bin` directory.

The command syntax is as follows:

```
– Windows restoreConfig.bat full_path_to_backup_file
```

```
– Linux UNIX ./restoreConfig.sh full_path_to_backup_file
```

For more information about the arguments of the **restoreConfig** command, see [WebSphere Application Server Knowledge Center](#).

d) Start the instances of the application server again.

- In a Network Deployment environment, complete the following steps:
  - a) Locate your backup configuration file.
 

The default directory is `dc_home/data`. If several backup files are present, check the modification date and time of the file; it must be the date and time of the failed configuration. If you did not complete any other data collector configurations on the same host after the failed one, use the most recent file in the directory.
  - b) Stop all instances of the application server.
  - c) Create a temporary directory in any convenient path (*temp\_directory*). On a UNIX or Linux system, create it under the `/tmp` directory.
  - d) Run the `restoreConfig` command from the `appserver_home/profiles/profile_name/bin` directory.
 

The command syntax is as follows:

    - **Windows** `restoreConfig.bat full_path_to_backup_file`
    - **Linux** | **UNIX** `./restoreConfig.sh full_path_to_backup_file`

The **restoreConfig** command restores the original application server configuration to the temporary directory.
  - e) Copy the `server.xml`, `variables.xml`, and `pmi-config.xml` files from temporary directory to the Deployment Manager system.
    - Source directory: `temp_directory/restored_configuration_home/cells/cell_name/nodes/node_name/servers/server_name`
    - Target directory: `appserver_home/profiles/profile_name/config/cells/cell_name/nodes/node_name/servers/server_name`
  - f) Complete a node sync from the Deployment Manager administrative console for the node.
  - g) In the Deployment Manager administrative console, save changes to the master configuration.
  - h) Start the instances of the application server.

## Configuring WebSphere Infrastructure Manager monitoring

Configure the WebSphere Infrastructure Manager agent to monitor the performance of WebSphere Deployment Manager and Node Agent.

### About this task

The WebSphere Infrastructure Manager agent is a multiple instance agent. You must create the first instance and start the agent manually.

### Procedure

1. To configure the agent, run the following command.

```
install_dir/bin/wim-agent.sh config instance_name
```

Where *instance\_name* is the name you want to give to the instance, and *install\_dir* is the installation directory of WebSphere Infrastructure Manager agent. The default installation directory is `/opt/ibm/apm/agent`.

2. When prompted to Edit 'Monitoring Agent for WebSphere Infrastructure Manager' settings, enter 1 to continue.
3. When prompted for Java home, specify the directory where Java is installed.
 

The default value is `/opt/ibm/apm/agent/JRE/1x8266/jre`.
4. When prompted for DMGR Profile Home, specify the home directory of the Deployment Manager profile.
 

The default directory is `/opt/IBM/WebSphere/AppServer/profiles/Dmgr01`.

5. When prompted for `JMX user ID`, specify the user ID that is used to connect to the MBean server.
6. When prompted to `Enter JMX password`, specify the password for the user.
7. When prompted to `Re-type JMX password`, enter the password again.
8. To start the agent, run the following command.

```
install_dir/bin/wim-agent.sh start instance_name
```

### What to do next

Log in to the Cloud App Management console to view monitoring data. For more information, see [“Starting the Cloud App Management UI”](#) on page 176.

## Configuring WebSphere MQ monitoring

Before you can start the IBM MQ(formerly WebSphere MQ) agent, you must assign an instance name to the agent and complete the several configuration tasks for the user ID.

### Before you begin

- The directions here are for the most current release of this agent. For information about how to check the version of an agent in your environment, see [Agent version command](#). For detailed information about the agent version list and what's new for each version, see [“Change history”](#) on page 52.
- Make sure that the system requirements for the IBM MQ(formerly WebSphere MQ) agent are met in your environment. For the up-to-date system requirement information, see the [Detailed system requirements report for the MQ agent](#).

### Procedure

1. Authorize the user ID that is used to configure, start, and stop the agent to access IBM MQ (WebSphere MQ) objects. See [“Authorizing the user IDs to run the agent”](#) on page 548.
2. Configure IBM MQ (WebSphere MQ) to enable the data that you want to monitor. See [“Configuring IBM MQ \(WebSphere MQ\) for data enablement”](#) on page 550.
3. Configure the agent by providing an agent instance name, a queue manager name, and optionally an agent name. See [“Configuring the IBM MQ\(formerly WebSphere MQ\) agent”](#) on page 551.

## Authorizing the user IDs to run the agent

For a user ID to configure, start, and stop the IBM MQ(formerly WebSphere MQ) agent, the user ID must belong to the **mqm** group, which has full administrative privileges over IBM MQ (WebSphere MQ). Also, for a non-root user or a non-administrator user, you must grant users the access to the IBM MQ (WebSphere MQ) objects by using the IBM MQ (WebSphere MQ) control command.

### About this task

On AIX or Linux system, you must add the user ID to the **mqm** group and then grant the user ID appropriate access to the IBM MQ (WebSphere MQ) objects with the **setmqaut** command.

On Windows systems, you must add the user ID to the **mqm** group. If the user ID does not belong to the Administrator user group, you must also use the Registry Editor to grant permissions to the user ID to start or stop the agent.

### Procedure

- |       |      |
|-------|------|
| Linux | UNIX |
|-------|------|

On AIX or Linux system, complete the following steps:

  - a) Log on to the AIX or Linux system by using the root ID.
  - b) Add the user ID that is used to run the agent to the **mqm** group.



- c) (WebSphere MQ V7.5 or later): If the user ID is a non-root user on the AIX or Linux system, set the appropriate level of authority for the user ID to access the IBM MQ (WebSphere MQ) objects by running the following command:

```
setmqaut -m queue_manager -t qmgr -p user_ID +inq +connect +dsp +setid
```

where *queue\_manager* is the name of the queue manager of WebSphere MQ V7.5 or later and *user\_ID* is the non-root or non-administrator user ID to run the agent.

- **Windows**

On Windows systems, complete the following steps:

- Log on to the Windows systems as a system administrator.
- Add the user ID that is used to run the agent to the **mqm** group.
- If the user ID that you use to start, run, and stop the agent is not a member of the Administrators group, use the Registry Editor to set permissions for a user ID to ensure that the agent can be started and stopped successfully:
  - Click **Start** > **Run**, and then type `regedit.exe` to open the Registry Editor.
  - In the Registry Editor, locate the key, `HKEY_LOCAL_MACHINE\SOFTWARE\Candle`.
  - Right-click the key and click **Permissions**.
  - If the user ID for the IBM MQ(formerly WebSphere MQ) agent is not in the Group or user names list, click **Add** to add the user ID to the list.
  - Click the user ID in the list.
  - In the Permissions for the *user-ID* list, where *user-ID* is the user ID of IBM MQ(formerly WebSphere MQ) agent, select **Full Control** in the Allow column and click **OK**.
  - In the Registry Editor, locate the key, `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib`.
  - Right-click the key and click **Permissions**.
  - If the user ID for the IBM MQ(formerly WebSphere MQ) agent is not in the Group or user names list, click **Add** to add the user ID to the list.
  - Click the user ID in the Group or user names list.
  - In the Permissions for the *user-ID* list, where *user-ID* is the user ID of IBM MQ(formerly WebSphere MQ) agent, select **Read** in the Allow column and click **OK**.
  - Close the Registry Editor.
- Locate the *install\_dir* directory, where *install\_dir* is the agent installation directory.
- Right-click the directory and click **Properties**.
- On the Security tab, if the user ID for IBM MQ(formerly WebSphere MQ) agent is not in the Group or user names list, click **Edit** and then **Add** to add the user ID to the list.
- Click the user ID in the Group or user names list.
- In the Permissions for the *user-ID* list, select **Full Control** in the Allow column, where **user-ID** is the user ID of IBM MQ(formerly WebSphere MQ) agent.
- Click **OK**.

### What to do next

The next step is to configure IBM MQ (WebSphere MQ) for data enablement. See [“Configuring IBM MQ \(WebSphere MQ\) for data enablement”](#) on page 550.

## Configuring IBM MQ (WebSphere MQ) for data enablement

Before you configure the IBM MQ (formerly WebSphere MQ) agent, it is recommended to configure the IBM MQ (WebSphere MQ) first to enable the data that you want to monitor.

### About this task

Decide what type of data that you want the IBM MQ (formerly WebSphere MQ) agent to monitor. Enable the data at the queue manager by using the MQSC commands if the data is not produced by the queue manager by default.

**Remember:** You must start MQSC for the target queue manager before you issue the MQSC commands. To get a list of the queue manager, issue the **dspmqr** command from the `bin` directory within the IBM MQ (WebSphere MQ) installation directory. To start MQSC for a queue manager, issue the following command from the `bin` directory, where `<qmgr_name>` is the name of the queue manager that you want to configure.

```
runmqsc <qmgr_name>
```

### Procedure

- To see the age of the oldest message on a queue, complete the steps as documented in [“Enabling real-time monitoring for queues”](#) on page 550.
- To monitor certain queue manager events that are not generated by the queue manager by default, complete the steps as documented in [“Enabling event monitoring for the queue manager”](#) on page 550.

### Enabling real-time monitoring for queues

#### About this task

To see the age of the oldest message (in seconds) on a queue, you must enable the real-time monitoring for the queue.

#### Procedure

Use the following commands to enable real-time monitoring for the queues in your environment.

- To enable real-time monitoring for all the queues whose MONQ attribute is set to QMGR, issue the following command:

```
ALTER QMGR MONQ(collection_level)
```

where *collection\_level* specifies the collection level of monitoring data for the queues. You can set it to LOW, MEDIUM, or HIGH to suit the requirements of your environment.

- To enable real-time monitoring for individual queue, issue the following command:

```
ALTER QLOCAL(queue_name) MONQ(collection_level)
```

where *queue\_name* is the name of the queue; *collection\_level* specifies the collection level of monitoring data for the queues. You can set it to LOW, MEDIUM, or HIGH to suit the requirements of your environment.

### Enabling event monitoring for the queue manager

#### About this task

Event monitoring is one of the monitoring techniques that are available to monitor your IBM MQ network. After you enable the queue manager to emit certain types of events, event messages are put on event queues when the event occurs. So that these event messages can be monitored and displayed by the IBM MQ (formerly WebSphere MQ) agent.

The following types of events are not monitored and displayed with the default queue manager configuration. Use the **ALTER QMGR** command to enable the queue manager to generate these events so that they can be displayed on the Cloud App Management console.

- Channel events
- Performance events

### Procedure

Use the following commands to enable the queue manager to generate the events that you care:

- To generate channel events, issue `ALTER QMGR CHLEV(ENABLED)`.
- To generate performance events, issue `ALTER QMGR PERFMEV(ENABLED)`.

## Configuring the IBM MQ(formerly WebSphere MQ) agent

You must assign an instance name to the IBM MQ(formerly WebSphere MQ) agent and configure the agent before it can start monitoring your IBM MQ (WebSphere MQ) environment.

### Before you begin

- Make sure that the agent user ID has appropriate permission to access IBM MQ (WebSphere MQ) objects. If you have not done it, follow the instructions in [“Authorizing the user IDs to run the agent”](#) on page 548.
- Configure IBM MQ (WebSphere MQ) to enable the required data collection. If you have not done it, see [“Configuring IBM MQ \(WebSphere MQ\) for data enablement”](#) on page 550.
- You must provide the name of queue manager to be monitored by the IBM MQ(formerly WebSphere MQ) agent. Contact the IBM MQ (WebSphere MQ) administrator if you do not know the appropriate queue manager name. Alternatively, issue the `dspsmq` command from the `bin` directory within the IBM MQ (WebSphere MQ) installation directory to get a list of the queue managers. The returned `QMNAME` value is what you must provide when you configure the IBM MQ(formerly WebSphere MQ) agent.

### About this task

The IBM MQ(formerly WebSphere MQ) agent is a multiple instance agent; you must create the first instance and manually start the agent.

On UNIX or Linux systems, you can choose to configure the agent with or without interactions. On Windows systems, you can configure the agent without interactions only.

- To configure the agent with interaction, run the configuration script and respond to prompts. See [“Interactive configuration”](#) on page 551.
- To configure the agent without interaction, edit the silent response file and then run the configuration script. See [“Silent configuration”](#) on page 552.

### Interactive configuration

#### Procedure

To configure the agent by running the script and responding to prompts, complete the following steps:

1. Enter the following command to create an agent instance:

```
install_dir/bin/mq-agent.sh config instance_name
```

where `install_dir` is the agent installation directory; `instance_name` is the name you want to give to the instance.

2. When prompted for Queue Manager Name, specify the name of the queue manager to be monitored.
3. When prompted for Agent Name, specify the agent name. Do not press Enter to skip specifying this parameter.

**Remember:** This agent name is different from the agent instance name. The agent instance name is used in the agent configuration file name to distinguish the configuration files between agents, for example, `hostname_mq_instancename.cfg`.

- When prompted for WebSphere MQ library path, press Enter to accept the default value, which is the 64-bit library path of IBM MQ (WebSphere MQ) automatically discovered by the IBM MQ (formerly WebSphere MQ) agent. If no default value is displayed, you must provide the 64-bit library path of IBM MQ (WebSphere MQ) to proceed.  
An example of the 64-bit library path is `/opt/mqm8/lib64` for a Linux system.

- To start the agent, enter the following command:

```
install_dir/bin/mq-agent.sh start instance_name
```

## Silent configuration

### Procedure

To configure the agent by editing the silent response file and running the script without interaction, complete the following steps:

- Open the `mq_silent_config.txt` file in a text editor.

- Linux** | **UNIX** `install_dir/samples/mq_silent_config.txt`
- Windows** `install_dir\tmaitm6_x64\samples\mq_silent_config.txt`

where `install_dir` is the agent installation directory.

- Required: For **QMNAME**, specify the name of the queue manager to be monitored.
- Required: For **AGTNAME**, specify an agent name.

**Remember:** This agent name is different from the agent instance name. The agent instance name is used in the agent configuration file name to distinguish the configuration files between agents, for example, `hostname_mq_instancename.cfg`.

- Optional: For **WMQLIBPATH**, specify the 64-bit library path of IBM MQ (WebSphere MQ). For example, `/opt/mqm8/lib64`. If no value is specified, the path can be automatically discovered during agent configuration.
- Save and close the `mq_silent_config.txt` file, and then run the following command from the command line:

- Linux** | **UNIX** `install_dir/bin/mq-agent.sh config instance_name path_to_responsefile`
- Windows** `install_dir\BIN\mq-agent.bat config instance_name path_to_responsefile`

where `instance_name` is the name of the instance that you configure, and `path_to_responsefile` is the full path of the silent response file.

**Remember:** On Windows systems, do not omit the double quotation marks (") that enclose the path to the silent response file, especially when the path contains special characters.

For example, if the response file is in the default directory, run the following command.

- Linux** | **UNIX**

```
/opt/ibm/apm/agent/bin/mq-agent.sh config instance_name /opt/ibm/apm/agent/samples/mq_silent_config.txt
```

- Windows**

```
C:\IBM\APM\BIN\mq-agent.bat config instance_name "C:\IBM\APM\tmaitm6_x64\samples\mq_silent_config.txt"
```

- To start the agent, enter the following command:

- **Linux** | **UNIX**

```
install_dir/bin/mq-agent.sh start instance_name
```

- **Windows**

```
install_dir\bin\mq-agent.bat start instance_name
```

## Results

Now, you can log in to the Cloud App Management user interface to view monitoring data. For more information, see [“Starting the Cloud App Management UI” on page 176](#).



---

## Chapter 15. Deploying ICAM Data Collectors

IBM Cloud App Management, Advanced provides two categories of ICAM Data Collectors to monitor your application workloads: Kubernetes data collector and runtime data collectors.

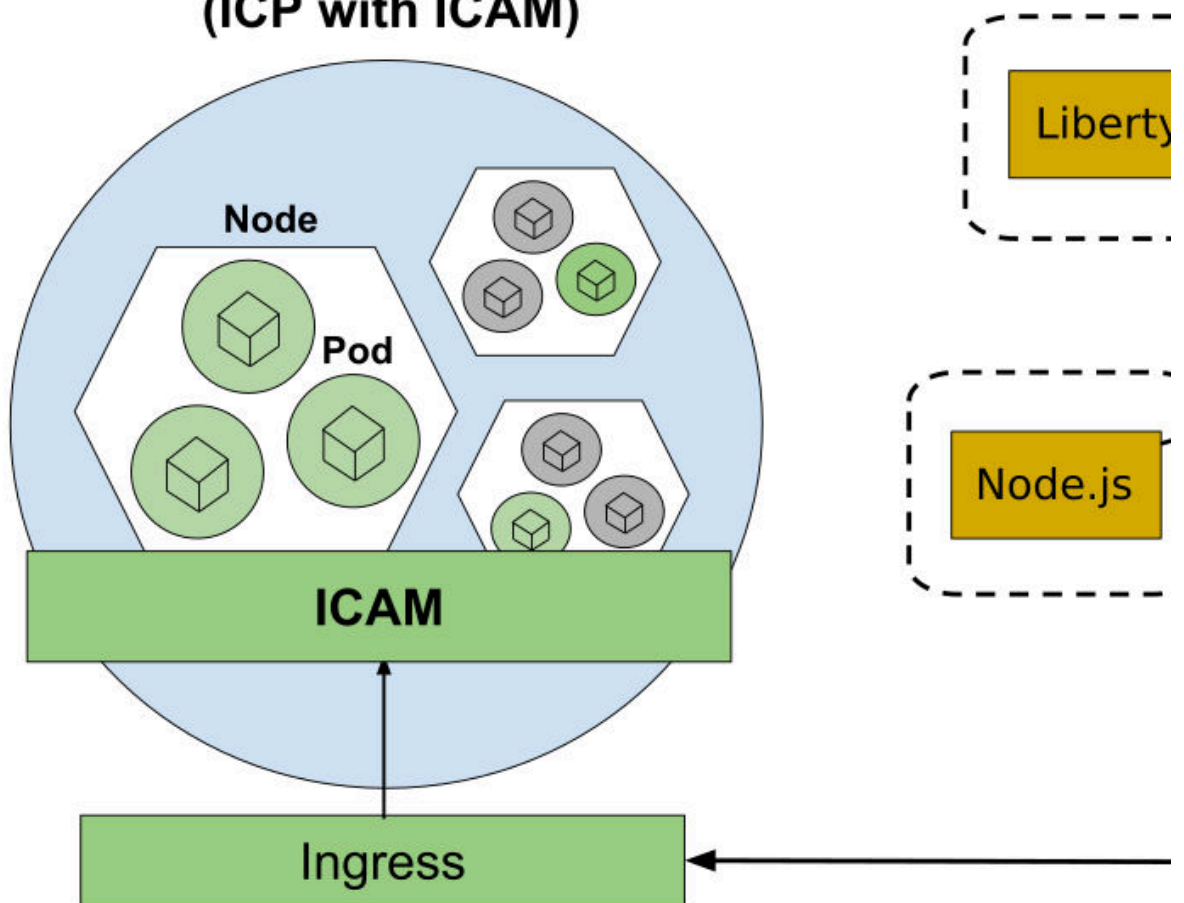
- The Kubernetes data collector is not embedded into the application workloads: it is installed outside the container workloads and typically collects the performance metrics by connecting to the workload's management interface. Kubernetes workloads can be monitored by the Kubernetes data collector.
- Runtime data collectors are embedded within the application workloads that you want to monitor, and can collect deep performance data about the application. Node.js, Liberty, and J2SE applications can be instrumented by runtime data collectors.

After downloading the data collector eImage from Passport Advantage and the configuration package from the Cloud App Management server, you can install and configure the Kubernetes data collector and runtime data collectors.

The illustration depicts an installation with the Kubernetes data collector installed on two clusters in the Kubernetes environment and the Liberty data collector and Node.js data collector installed on one cluster.

---

## Kubernetes Cluster (ICP with ICAM)



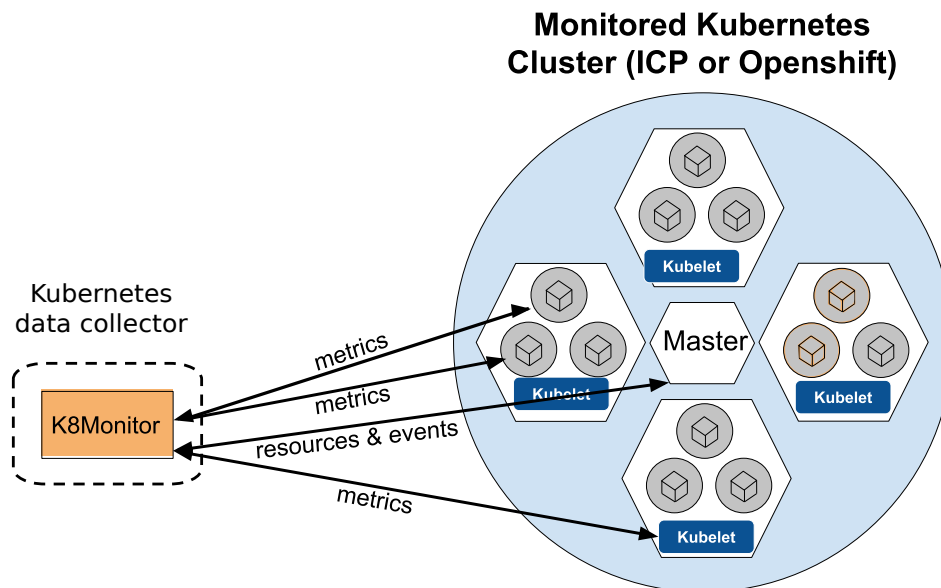


## Kubernetes data collector

The Kubernetes data collector manages the collection, enrichment, and dispatch of Kubernetes topology, event, and performance data. You can install the data collector directly on your IBM Cloud App Management cluster, on a remote cluster, or both.

The data collector is installed on each Kubernetes cluster that you want to monitor. You can deploy the Kubernetes data collector on OpenShift and IBM Cloud Private platforms either by running an Ansible script or by entering the commands manually.

The illustration shows the data flow of metrics and events from the Kubernetes workloads to the K8Monitor component.



## Configuring Kubernetes monitoring

After installing your Cloud App Management server, you can configure the Kubernetes data collector for monitoring the applications in your Kubernetes environment. Use the Kubernetes data collector to manage the collection, enrichment, and dispatch of Kubernetes topology, events, and performance data.

### Before you begin

Prerequisites:

- Ansible version 2.4.2 or higher. Otherwise, follow the instructions in [“Configuring Kubernetes monitoring without Ansible”](#) on page 562
- Helm client and server (Tiller) version 2.11.0 or higher. Otherwise, follow the instructions in [“Configuring Kubernetes monitoring without Helm”](#) on page 566
- Kubectl client on the environment from where you are installing
- Kubernetes version 1.7 or higher (available with IBM Cloud Private version 3.2.0 or higher and OpenShift version 3.9 or higher)

Sizing requirements:

- As described in [“Planning hardware and sizing”](#) on page 77, your monitored environment can be either Size0 or Size1. For Kubernetes monitoring, Size0 is for monitoring up to 1000 containers and requires 1

GB RAM and 0.5 CPU cores, Size1 is for monitoring up to 4000 containers and requires 2 GB RAM and 2.0 CPU cores.

#### Considerations:

- If you want to deploy a Kubernetes data collector that is configured to point to another Cloud App Management server on the same cluster, you must deploy it in a different namespace so as not to disrupt the configuration secret (or secrets) in use by active releases.
- If you are installing the Kubernetes data collector into another namespace, you must also assign `docker_group` and with a value that matches the namespace. Example: `ansible-playbook helm-main.yaml --extra-vars="cluster_name=myCluster50 release_name=camserver namespace=sample docker_group=sample tls_enabled=true"`
- If you are installing on an OpenShift environment, you might first need to grant the Helm Tiller service edit access to the project or namespace where you want to install the Kubernetes data collector. For more information, see <https://blog.openshift.com/getting-started-helm-openshift/>.
- When you install on an OpenShift environment, you might need to override the default security configuration. Otherwise, it is possible for the user ID to be at variance with what is expected by the application image, resulting in exceptions such as permission errors. For more information, see the OpenShift Cookbook topic, [How can I enable an image to run as a set user ID?](#)
- If your environment configuration requires adding entries to a Pod's `/etc/hosts` file (to provide Pod-level overrides of hostname resolution when DNS and other options are not applicable), specify the appropriate `hostAliases` in this helm chart's `values.yaml` file. For more information on `hostAliases`, please see the [Kubernetes documentation](#).

#### Connectivity check:

- Before installing the Kubernetes data collector, ensure that the ingresses to the Cloud App Management server are accessible by running the healthcheck command. Port 443 is the SSL (Single Socket Layer) connection for the Cloud App Management browser client, and is also used by monitoring agents that connect to the server using SSL. Port 80 is the HTTP connection for agents that are not configured for SSL.

```
curl -k https://my_ICP_proxy_node_IP:443/applicationmgmt/1.0/healthcheck
curl -k http://my_ICP_proxy_node_IP:80/applicationmgmt/1.0/healthcheck
```

where `my_ICP_proxy_node_IP` is the IP address of the IBM Cloud Private proxy node. (such as `curl -k https://icam_a1_raleigh.ibm.com:443/applicationmgmt/1.0/healthcheck`). The ports are the incoming connections that are going to use SSL (port 443) or non-SSL (port 80).

### About this task

Deploying the Kubernetes data collector involves downloading the data collectors installation eImage, logging into the Cloud App Management console and downloading the data collector configuration package, installing the data collector, and validating the installation.

The eImage is the data collectors package and contains all the installable data collectors. The configuration package (ConfigPack) contains the ingress URLs and authentication information required to configure the data collector package to communicate with the Cloud App Management server.

### Procedure

Download the eImage data collectors installation tar file and the data collector configuration package:

1. If you haven't already, download the data collectors installation eImage (part number CC5H0EN) from [IBM Passport Advantage](#).

For more information, see [“Part numbers” on page 71](#).

2. Download the data collector configuration package:

- a) Log in to the Cloud App Management console, click the **Get Started** link on the Welcome page, then select **Administration > Integrations**.
- b) Click the **New integration** button.

- c) In the **Standard monitoring agents** section, select the **ICAM Data Collectors Configure** button.
  - d) Select **Download file** and specify the directory where you want to save the compressed data collector configuration package, `ibm-cloud-apm-dc-configpack.tar`.
3. Move the downloaded installation package and the configuration package to a node in the cluster that you want to monitor:

Examples using secure copy:

```
scp my_path_to_download/app_mgmt_k8sdc.tar.gz
root@my.env.com:/my_path_to_destination
scp my_path_to_download/ibm-cloud-apm-dc-configpack.tar root@my.env.com:
/my_path_to_destination
```

where

`my_path_to_download` is the path to where the installation tar file or configuration package file was downloaded


`root@my.env.com` is your user ID on the system where the kubectl client is configured to point to the environment to be monitored

`my_path_to_destination` is the path to the environment that you want to monitor

Install the Kubernetes data collector in the Kubernetes cluster that you want to monitor:

4. If you are not installing from your master node, configure the kubectl client to point to the master node of the cluster that you want to monitor.

This step isn't needed if you are installing from your master node because the kubectl client points to the node that you are on by default.

In the IBM Cloud Private management console, you can click  > **Configure client** and follow the instructions to run the **kubectl config** commands.

5. Initialize Helm:

```
helm init
```

6. Log in to your Docker registry.

The Docker registry must be the same one referenced in the Ansible script command.

```
docker login -u my_username -p my_password my_clustername:my_clusterport
```

where

`my_username` and `my_password` are the user name and password for the Docker registry

`my_clustername` is the name of the cluster that you're monitoring

`my_clusterport` is the port number for the Docker registry

7. Extract the Kubernetes data collector package from the installation tar file that you downloaded in step “3” on page 559 and move `ibm-cloud-apm-dc-configpack.tar` to your working directory:

```
tar -xvf appMgtDataCollectors_2019.4.0.2.tar.gz
cd appMgtDataCollectors_2019.4.0.2
tar -xvf app_mgmt_k8sdc.tar.gz
cd app_mgmt_k8sdc
mv my_path_to_configpack/ibm-cloud-apm-dc-configpack.tar
```

8. If you want to use the default environment size, which is `size0`, you can skip this step. If you want to use another deployment size, complete this step.

Edit the `values.yaml` file to specify `environmentSize`:

```
tar -xvf app_mgmt_k8sdc_helm.tar.gz --warning=no-timestamp
sed -i 's/environmentSize:.*environmentSize: "size1"/' k8monitor/values.yaml
mv app_mgmt_k8sdc_helm.tar.gz app_mgmt_k8sdc_helm_old.tar.gz
tar -cvf app_mgmt_k8sdc_helm.tar k8monitor/
gzip app_mgmt_k8sdc_helm.tar
```

where `size1` is the environment size (see “[Planning hardware and sizing](#)” on page 77).

Before running the installation script in the next step, you must decide on important deployment configurations (Table 78 on page 560) to be passed to the install script.

- Use Table 78 on page 560 to determine which options to specify, then run the Ansible script with the configuration options and defaults that are required for the Kubernetes environment that you are monitoring:

```
ansible-playbook helm-main.yaml --extra-vars="configOption1=configValue1
configOption2=configValue2"
```

Examples:

```
ansible-playbook helm-main.yaml --extra-vars="cluster_name=myCluster
release_name=camserver namespace=default docker_group=default tls_enabled=true"
```

Table 78. Ansible-playbook configuration options

Configuration option	Description	Required	Default configuration value
cluster_name	<p>Unique name to distinguish your cluster from the other clusters being monitored. Only alphanumeric characters and "-" are supported, with no spaces. If you enter invalid characters, they are removed from the name. The assigned cluster name in the example is myCluster. It is not recommended that you change the cluster name after deployment. If you must change the cluster_name after deployment, see "What to do next" later in this topic and be advised that:</p> <ol style="list-style-type: none"> <li>The change will not be immediately available. You must wait for the monitor to restart and complete one collection cycle.</li> <li>Previous thresholds and policies that denote the cluster may need to be manually updated.</li> <li>Preexisting incidents will reference the old cluster name until they expire.</li> </ol>	No	UnnamedCluster
release_name	The Helm release name for the Kubernetes data collector. Choose a release name that does not yet exist in your environment.	No	icam-kubernetes-resources
namespace	The Kubernetes namespace where you want your Kubernetes data collector and configuration secrets to be created. The namespace must already exist in your cluster, because it will not be created.	No	default
docker_registry	The host and port of the Docker registry where you want to store the data collector images. Example: mycluster.icp:8500.	No	mycluster.icp:8500

Configuration option	Description	Required	Default configuration value
docker_group	The Docker group in the registry where you want to store your images. Example: <code>myRegistry:1000/mydockergroup</code> . If you are installing the data collector in a different namespace from the default, you must also assign <code>docker_group</code> and with the same name as the namespace.	No	default
tls_enabled	Specifies whether TLS (Transport Layer Security) is enabled in your environment: <code>true</code> or <code>false</code> .	Yes	No default. You must provide a value: <code>true</code> or <code>false</code>




**Trouble:** In a slower environment, the Ansible script might fail during the TASK [Gather Facts] stage due to a time out and return a message such as "Timer expired after 10 seconds". If this happens, edit the Ansible config file at `/etc/ansible/ansible.cfg` by uncommenting the `# gather_timeout = 10` line and extending the time out value (30 should be sufficient).

Validate the deployment:

- After the installation script has completed, wait for the deployment to become ready as indicated by this message:

```
kubectl get deployment my_ReleaseName-k8monitor --namespace=myReleaseNamespace
```

Depending on the size and health of your environment, it can take up to 10 minutes for the Kubernetes data collector to start up and output logs that you can review (see [“Checking the Kubernetes installation logs”](#) on page 569). The data collector startup creates a Kubernetes event, which generates an informational incident.


- View the data collector metrics and incidents in the Cloud App Management console to confirm that the data collector is successfully monitoring:
  - Select the **Resources** tab. Find your Kubernetes resource types. For instructions, see [“Viewing your managed resources”](#) on page 769. If this is your first installation, you'll see 1 cluster.
  - Select the **Incidents** tab and click **All incidents**, then click  and filter by **Priority 4** incidents. You should see incidents about Kubernetes monitoring availability. For more information, see [“Managing incidents”](#) on page 747.

You can also review the logs as described in [“Checking the Kubernetes installation logs”](#) on page 569.

## Results

- The Kubernetes data collector is installed and begins sending metrics to the Cloud App Management server for display in the **Resource** dashboard pages. Incidents are generated for any native Kubernetes events.
- The `ibm-k8monitor-config` ConfigMap is created in your default namespace as part of Kubernetes data collector deployment. The ConfigMap contains the `ProviderId` that is used to distinguish this cluster's resources from the others in your tenant namespace and is crucial to enable multi-cluster support. Do not delete, move, or rename this resource. If this ConfigMap is deleted and the Kubernetes data collector is restarted or is deployed or redeployed, your data duplicates itself within the tenant because the monitor sees it as a new cluster.

## What to do next

- For each Kubernetes cluster that you want to monitor, repeat the steps starting at step [“3” on page 559](#) to install the Kubernetes data collector and validate the deployment.
- If you reconfigure or provide your own ingress certificates post-deployment, you must restart the agent bootstrap service, download the updated ConfigPack, and reconfigure your deployed data collectors to use the updated configurations. For more information, see [“Configuring a custom server certificate” on page 153](#).
- If you want to change your `cluster_name` post-deployment, enter the following command and change the `CLUSTER_NAME` environment variable: `kubectl edit deployment my_ReleaseName-k8monitor`.
- For troubleshooting the deployment, see [“Kubernetes data collector issues” on page 575](#).
- Use the resource dashboards and create thresholds to monitor your Kubernetes environment. For more information, see [“Viewing your managed resources” on page 769](#) and [“Kubernetes metrics for thresholds” on page 570](#).
- To learn more about Kubernetes and best practices in managing your environment, see the PDF on IBM Cloud App Management Developer Center: [Kubernetes, CPU, and memory management](#) .

## Configuring Kubernetes monitoring without Ansible

After installing your Cloud App Management server, you can configure the Kubernetes data collector for monitoring the applications in your Kubernetes environment. Use this procedure if you have no Ansible server. The Kubernetes data collector manages the collection, enrichment, and dispatch of Kubernetes topology, event, and performance data.

### Before you begin

Prerequisites:

- Helm client and server (Tiller) version 2.11.0 or higher. Otherwise, follow the instructions in [“Configuring Kubernetes monitoring without Helm” on page 566](#)
- Kubectl client on the environment from where you are installing
- Kubernetes version 1.7 or higher (available with IBM Cloud Private version 3.2.0 or higher and OpenShift version 3.9 or higher)

Considerations:

- If you want to deploy a Kubernetes data collector that is configured to point to another Cloud App Management server on the same cluster, you must deploy it in a different namespace so as not to disrupt the configuration secret (or secrets) in use by active releases.
- If you are installing the Kubernetes data collector into another namespace, you must also assign `docker_group` and with a value that matches the namespace. Example: `ansible-playbook helm-main.yaml --extra-vars="cluster_name=myCluster50 release_name=camservice namespace=sample docker_group=sample tls_enabled=true"`
- If you are installing on an OpenShift environment, you might first need to grant the Helm Tiller service edit access to the project or namespace where you want to install the Kubernetes data collector. For more information, see <https://blog.openshift.com/getting-started-helm-openshift/>.
- When you install on an OpenShift environment, you might need to override the default security configuration. Otherwise, it is possible for the user ID to be at variance with what is expected by the application image, resulting in exceptions such as permission errors. For more information, see the OpenShift Cookbook topic, [How can I enable an image to run as a set user ID?](#)
- If your environment configuration requires adding entries to a Pod's `/etc/hosts` file (to provide Pod-level overrides of hostname resolution when DNS and other options are not applicable), specify the appropriate `hostAliases` in this helm chart's `values.yaml` file. For more information on `hostAliases`, please see the [Kubernetes documentation](#).

## About this task

Deploying the Kubernetes data collector involves downloading the data collectors installation eImage, logging into the Cloud App Management console and downloading the data collector configuration package, installing the data collector, and validating the installation.

The eImage is the data collectors package and contains all the installable data collectors. The configuration package (ConfigPack) contains the ingress URLs and authentication information required to configure the data collector package to communicate with the Cloud App Management server.

## Procedure

Download the eImage data collectors installation tar file and the data collector configuration package:

1. If you haven't already, download the data collectors installation eImage (part number CC5H0EN) from [IBM Passport Advantage](#).

For more information, see [“Part numbers” on page 71](#).

2. Download the data collector configuration package:
  - a) Log in to the Cloud App Management console, click the **Get Started** link on the Welcome page, then select **Administration > Integrations**.
  - b) Click the **New integration** button.
  - c) In the **Standard monitoring agents** section, select the **ICAM Data Collectors Configure** button.
  - d) Select **Download file** and specify the directory where you want to save the compressed data collector configuration package, `ibm-cloud-apm-dc-configpack.tar`.
3. Move the downloaded installation package and the configuration package to a node in the cluster that you want to monitor:

Examples using secure copy:

```
scp my_path_to_download/app_mgmt_k8sdc.tar.gz
root@my.env.com:/my_path_to_destination
scp my_path_to_download/ibm-cloud-apm-dc-configpack.tar root@my.env.com:
/my_path_to_destination
```

where


`my_path_to_download` is the path to where the installation tar file or configuration package file was downloaded

`root@my.env.com` is your user ID on the system where the kubectl client is configured to point to the environment to be monitored

`my_path_to_destination` is the path to the environment that you want to monitor

Install the Kubernetes data collector in the Kubernetes cluster that you want to monitor:

4. If you are not installing from your master node, configure the kubectl client to point to the master node of the cluster that you want to monitor.

In the IBM Cloud Private management console, you can click  > **Configure client** and follow the instructions to run the **kubectl config** commands.

5. Initialize Helm:

```
helm init
```

6. Log in to your Docker registry.

```
docker login -u my_username -p my_password my_clustername:my_clusterport
```

where

`my_username` and `my_password` are the user name and password for the Docker registry

`my_clustername` is the name of the cluster that you're monitoring

`my_clusterport` is the port number for the Docker registry

7. Extract the Kubernetes data collector installation package from the installation tar file that you downloaded in step 3:

```
tar -xvf appMgtDataCollectors_2019.4.0.2.tar.gz
cd appMgtDataCollectors_2019.4.0.2
tar -xvf app_mgmt_k8sdc.tar.gz
cd app_mgmt_k8sdc
```

8. Extract the data collector configuration package file that you secure copied in step 3:

```
tar -xvf my_path_to/ibm-cloud-apm-dc-configpack.tar
```

The data collector ConfigPack is extracted to the `appMgtDataCollectors_2019.4.0` directory.

9. Load the Docker images:

```
docker load -i app_mgmt_k8sdc_docker.tar.gz
```

10. Discover the k8-monitor Docker image repositories and tags:

```
K8_MONITOR_IMAGE_REPO=`docker images | grep icam-k8-monitor | head -1 | awk
'{{print $1}}'`
K8_MONITOR_IMAGE_TAG=`docker images | grep icam-k8-monitor | grep APM | head -1
| awk '{{print $2}}'`
```

11. Create the required configuration and security secrets:

```
kubectl -n my_namespace create -f ibm-cloud-apm-dc-configpack/dc-secret.yaml
kubectl -n my_namespace create secret generic ibm-agent-https-secret \
--from-file=ibm-cloud-apm-dc-configpack/keyfiles/cert.pem \
--from-file=ibm-cloud-apm-dc-configpack/keyfiles/ca.pem \
--from-file=ibm-cloud-apm-dc-configpack/keyfiles/key.pem
```

where

`my_namespace` is the namespace where you want your configuration secrets to be created.

12. Tag and push the Docker images to the Docker registry:

```
docker tag $K8_MONITOR_IMAGE_REPO:$K8_MONITOR_IMAGE_TAG my_docker_registry:
my_docker_registry_port/my_docker_group/k8-monitor:$K8_MONITOR_IMAGE_TAG
docker push my_docker_registry:my_docker_registry_port/my_docker_group
/k8-monitor:$K8_MONITOR_IMAGE_TAG
```

where

`my_docker_registry` is the host of the Docker registry where you want to store the image. Example: `mycluster.icp`.

`my_docker_registry_port` is the Docker registry service port. For example: 8500

`my_docker_group` is the Docker group in the registry where you want to store your images.

Example: `myRegistry:8500/mydockergroup`.

If you are installing the data collector in a different namespace from the default, you must also assign `docker_group` with the same name as the namespace.

13. Install the Helm Chart:

Install the Helm Chart with HTTPS enabled. If TLS is not enabled, do not include `--tls` in the last `set` command:

```
helm install app_mgmt_k8sdc_helm.tar.gz --name my_release_name --namespace my_namespace \
--set k8monitor.image.repository=my_docker_registry:my_docker_registry_port \
--set k8monitor.clusterName=my_cluster_name \
--set k8monitor.imageNamePrefix=my_docker_group/ \
--set k8monitor.imageTag=$K8_MONITOR_IMAGE_TAG \
--set k8monitor.ibmAgentConfigSecret=dc-secret \
--set k8monitor.ibmAgentHTTPSSecret=ibm-agent-https-secret
```

where



*my\_release\_name* is the Helm release name for data collector. Choose a release name that does not yet exist in your environment.

*my\_namespace* is the namespace where you want your data collector to be installed.

*my\_docker\_registry* is the host of the Docker registry where the image is stored.

*my\_docker\_registry\_port* is the Docker registry service port.

*my\_cluster\_name* is a unique name to distinguish your cluster from the other clusters being monitored. Only alphanumeric characters and "-" are supported, with no spaces. If you enter invalid characters, they are removed from the name. It is not recommended that you change the cluster name after deployment

*my\_docker\_group* is the Docker group in the registry where you want to store your images.

Example: `myRegistry:8500/mydockergroup`.


If you are installing the data collector in a different namespace from the default, you must also assign `docker_group` with the same name as the namespace.

Validate the deployment:

14. After the installation script has completed, wait for the deployment to become ready as indicated by this message:

```
kubectl get deployment my_ReleaseName-k8monitor --namespace=myReleaseNamespace
```

Depending on the size and health of your environment, it can take up to 10 minutes for the Kubernetes data collector to start up and output logs that you can review (see [“Checking the Kubernetes installation logs”](#) on page 569). The data collector startup creates a Kubernetes event, which generates an informational incident.

15. View the data collector metrics and incidents in the Cloud App Management console to confirm that the data collector is successfully monitoring:
  - Select the **Resources** tab. Find your Kubernetes resource types. For instructions, see [“Viewing your managed resources”](#) on page 769. If this is your first installation, you'll see 1 cluster.
  - Select the **Incidents** tab and click **All incidents**, then click  and filter by **Priority 4** incidents. You should see incidents about Kubernetes monitoring availability. For more information, see [“Managing incidents”](#) on page 747.


You can also review the logs as described in [“Checking the Kubernetes installation logs”](#) on page 569.

## Results

- The Kubernetes data collector is installed and begins sending metrics to the Cloud App Management server for display in the **Resource** dashboard pages. Incidents are generated for any native Kubernetes events.
- The `ibm-k8monitor-config` ConfigMap is created in your default namespace as part of Kubernetes data collector deployment. The ConfigMap contains the `ProviderId` that is used to distinguish this cluster's resources from the others in your tenant namespace and is crucial to enable multi-cluster support. Do not delete, move, or rename this resource. If this ConfigMap is deleted and the Kubernetes data collector is restarted or is deployed or redeployed, your data duplicates itself within the tenant because the monitor sees it as a new cluster.

## What to do next

- For each Kubernetes cluster that you want to monitor, repeat the steps starting with step 3 to install the Kubernetes data collector and validate the deployment.
- If you reconfigure or provide your own ingress certificates post-deployment, you must restart the agent bootstrap service, download the updated ConfigPack, and reconfigure your deployed data collectors to use the updated configurations. For more information, see [“Configuring a custom server certificate”](#) on page 153.

- If you want to change your `cluster_name` post-deployment, enter the following command and change the `CLUSTER_NAME` environment variable: `kubectl edit deployment my_ReleaseName-k8monitor`.
- For troubleshooting the deployment, see [“Kubernetes data collector issues” on page 575](#).
- Use the resource dashboards and create thresholds to monitor your Kubernetes environment. For more information, see [“Viewing your managed resources” on page 769](#) and [“Kubernetes metrics for thresholds” on page 570](#).
- To learn more about Kubernetes and best practices in managing your environment, see the PDF on IBM Cloud App Management Developer Center: [Kubernetes, CPU, and memory management](#) .

## Configuring Kubernetes monitoring without Helm

After installing your Cloud App Management server, you can configure the Kubernetes data collector for monitoring the applications in your Kubernetes environment. This procedure is for environments with no Helm installation, such as OpenShift. The Kubernetes data collector manages the collection, enrichment, and dispatch of Kubernetes topology, event, and performance data.

### Before you begin

Prerequisites:

- Kubectl client on the environment from where you are installing
- Kubernetes version 1.7 or higher (available with IBM Cloud Private version 3.2.0 or higher and OpenShift version 3.9 or higher)

Considerations:

- If you are installing the Kubernetes data collector into another namespace, you must also assign `docker_group` and with a value that matches the namespace.
- When you install on an OpenShift environment, you might need to override the default security configuration. Otherwise, it is possible for the user ID to be at variance with what is expected by the application image, resulting in exceptions such as permission errors. For more information, see the OpenShift Cookbook topic, [How can I enable an image to run as a set user ID?](#)
- On the OpenShift platform, you can replace **kubectl** with the **oc** command.
- If your environment configuration requires adding entries to a Pod's `/etc/hosts` file (to provide Pod-level overrides of hostname resolution when DNS and other options are not applicable), specify the appropriate `hostAliases` in this helm chart's `values.yaml` file. For more information on `hostAliases`, please see the Kubernetes documentation <https://kubernetes.io/docs/concepts/services-networking/add-entries-to-pod-etc-hosts-with-host-aliases/>.

### About this task

Deploying the Kubernetes data collector involves downloading the data collectors installation eImage, logging into the Cloud App Management console and downloading the data collector configuration package, installing the data collector, and validating the installation.

The eImage is the data collectors package and contains all the installable data collectors. The configuration package (ConfigPack) contains the ingress URLs and authentication information required to configure the data collector package to communicate with the Cloud App Management server.

### Procedure

Download the eImage data collectors installation tar file and the data collector configuration package:

1. If you haven't already, download the data collectors installation eImage (part number CC5H0EN) from [IBM Passport Advantage](#).

For more information, see [“Part numbers” on page 71](#).

2. Extract the Docker images:

```
tar xvf appMgtDataCollectors_2019.4.0.2.tar.gz images/
```

3. Log in to your Docker registry.

```
docker login -u my_username -p my_password my_cluster_ca_domain:my_docker_registry_port
```

where

*my\_username* and *my\_password* are the user name and password for the Docker registry  
*my\_cluster\_ca\_domain* is the target cluster CA domain to monitor  
*my\_docker\_registry\_port* is the Docker registry service port. For example: 8500

4. Load and push the images to your Docker repository:

a) Load the k8-monitor Docker image to the repository:

```
docker load -i app_mgmt_k8sdc/app_mgmt_k8sdc_docker.tar.gz
docker tag my_repotags my_cluster_ca_domain:my_docker_registry_port/my_namespace/k8-
monitor:my_image_tag
docker push my_clustername:my_clusterport/my_namespace/k8-monitor:my_image_tag
```

b) Load the k8sdc-operator Docker image to the repository:

```
docker load -i app_mgmt_k8sdc/app_mgmt_k8sdc_operator.tar.gz
docker tag my_repotags my_cluster_ca_domain:my_docker_registry_port/my_namespace/k8sdc-
operator:my_image_tag
docker push my_cluster_ca_domain:8500/my_namespace/k8sdc-operator:my_image_tag
```

where:

*my\_username* and *my\_password* are the user name and password for the Docker registry  
*my\_cluster\_ca\_domain* is the target cluster CA domain to monitor  
*my\_docker\_registry\_port* is the Docker registry service port. For example: 8500  
*my\_namespace* is the target namespace on the cluster  
*my\_repotags* and *my\_image\_tag* are the Docker image tags, which you can get from the RepoTags in the manifest .son within the archive file.

5. Create Docker **imagePullSecrets**:

```
kubectl config set-context my_cluster_ca_domain-context --user=my_username --
namespace=my_namespace
kubectl create secret docker-registry my_registrykey \
--docker-server=my_cluster_ca_domain:my_docker_registry_port \
--docker-username=my_username \
--docker-password=my_password \
--docker-email=my_user_email
kubectl get secret
```

6. Download the data collector configuration package:

- Log in to the Cloud App Management console, click the **Get Started** link on the Welcome page, then select **Administration > Integrations**.
- Click the **New integration** button.
- In the **Standard monitoring agents** section, select the **ICAM Data Collectors Configure** button.
- Select **Download file** and specify the directory where you want to save the compressed data collector configuration package, `ibm-cloud-apm-dc-configpack.tar`.

7. Create the required configuration and security secrets:

```
kubectl -n my_namespace create -f ibm-cloud-apm-dc-configpack/dc-secret.yaml
kubectl -n my_namespace create secret generic ibm-agent-https-secret \
--from-file=ibm-cloud-apm-dc-configpack/keyfiles/cert.pem \
--from-file=ibm-cloud-apm-dc-configpack/keyfiles/ca.pem \
--from-file=ibm-cloud-apm-dc-configpack/keyfiles/key.pem
```

where *my\_namespace* is the target namespace on the cluster

8. Deploy the Kubernetes data collector resources. Apply the files under the `deploy` directory for defining and deploying the `k8sdc-operator`:

a) Create the "custom resource definition" for the operator spec:

```
kubectl create -f deploy/crds/k8sdc_crd.yaml
```

b) Apply any necessary value changes to the "Custom Resource" for operator spec by editing the `deploy/crds/k8sdc_cr.yaml` files

```
repository: "my_cluster_ca_domain:my_docker_registry_port"
imageNamePrefix: "my_namespace"
```

where `my_cluster_ca_domain` is the target cluster CA domain to monitor and `my_namespace` is the target namespace on the cluster

```
kubectl create -f deploy/crds/k8sdc_cr.yaml
```

c) Create the service-account for the operator:

```
kubectl create -f deploy/service_account.yaml
```

d) Apply the **imagePullSecrets** that you created in step "5" on page 567 to create the "k8sdc-operator" service account:

```
kubectl patch serviceaccount k8sdc-operator -p '{"imagePullSecrets": [{"name": "my_pull_secret_name"}]}'
```

**Note:** Bind cluster-admin with the "k8sdc-operator" service account. For OpenShift monitoring, you must create a **ClusterRoleBinding** using the following command:

```
oc create clusterrolebinding my_cluster_role_binding_name
--clusterrole=cluster-admin
--serviceaccount=my_namespace:k8sdc-operator -n my_namespace
```

where `my_cluster_role_binding_name` is a new cluster role binding name and `my_namespace` is the target namespace on the cluster (the project name in OpenShift).

e) Create the operator role:

```
kubectl create -f deploy/role.yaml
```

f) Create the operator role binding:

```
kubectl create -f deploy/role_binding.yaml
```

g) Create the `k8sdc-operator` by editing the `deploy/operator.yaml`:

**image:** `my_cluster_ca_domain:my_docker_registry_port/my_namespace/k8sdc-operator:my_imagetag`

```
kubectl create -f deploy/operator.yaml
```


Validate the deployment:

9. After the installation script has completed, wait for the deployment to become ready as indicated by this message:

```
kubectl get deployment my_ReleaseName-k8sdc-k8monitor --namespace=myReleaseNamespace
```

Depending on the size and health of your environment, it can take up to 10 minutes for the Kubernetes data collector to start up and output logs that you can review (see "[Checking the Kubernetes installation logs](#)" on page 569). The data collector startup creates a Kubernetes event, which generates an informational incident.

10. View the data collector metrics and incidents in the Cloud App Management console to confirm that the data collector is successfully monitoring:


- Select the **Resources** tab. Find your Kubernetes resource types. For instructions, see [“Viewing your managed resources” on page 769](#). If this is your first installation, you'll see 1 cluster.
- Select the **Incidents** tab and click **All incidents**, then click  and filter by **Priority 4** incidents. You should see incidents about Kubernetes monitoring availability. For more information, see [“Managing incidents” on page 747](#).

You can also review the logs as described in [“Checking the Kubernetes installation logs” on page 569](#).

## Results

- The Kubernetes data collector is installed and begins sending metrics to the Cloud App Management server for display in the **Resource** dashboard pages. Incidents are generated for any native Kubernetes events.
- The `ibm-k8monitor-config` ConfigMap is created in your default namespace as part of Kubernetes data collector deployment. The ConfigMap contains the `ProviderId` that is used to distinguish this cluster's resources from the others in your tenant namespace and is crucial to enable multi-cluster support. Do not delete, move, or rename this resource. If this ConfigMap is deleted and the Kubernetes data collector is restarted or is deployed or redeployed, your data duplicates itself within the tenant because the monitor sees it as a new cluster.

## What to do next

- For each Kubernetes cluster that you want to monitor, repeat the steps to install the Kubernetes data collector and validate the deployment.
- If you reconfigure or provide your own ingress certificates post-deployment, you must restart the agent bootstrap service, download the updated ConfigPack, and reconfigure your deployed data collectors to use the updated configurations. For more information, see [“Configuring a custom server certificate” on page 153](#).
- If you want to change your `cluster_name` post-deployment, enter the following command and change the `CLUSTER_NAME` environment variable: `kubectl edit deployment my_ReleaseName-k8monitor`.
- For troubleshooting the deployment, see [“Kubernetes data collector issues” on page 575](#).
- Use the resource dashboards and create thresholds to monitor your Kubernetes environment. For more information, see [“Viewing your managed resources” on page 769](#) and [“Kubernetes metrics for thresholds” on page 570](#).
- To learn more about Kubernetes and best practices in managing your environment, see the PDF on IBM Cloud App Management Developer Center: [Kubernetes, CPU, and memory management](#) .

## Checking the Kubernetes installation logs

Review the installation logs to ensure that the Kubernetes data collector is configured successfully.

### Procedure

1. Get the full name of your Kubernetes data collector by running the following command:

```
kubectl get pods --selector=app=k8monitor
```

In the following example, the full name of the Kubernetes data collector is `icamkubedc-k8monitor-7c4879844c-bxcgp`

2. Review the K8Monitor logs to ensure that all initial checks passed by running the following command:

```
kubectl logs -f full_kubernetes_data_collector_name-k8monitor --namespace=my_Namespace resource-event-collector
```

Look for key success messages:

```
- "ProviderId SOMEVALUEHERE successfully loaded from existing configMap"
```

```
(Or newly created)
- "Successfully initialized in-cluster kubernetes monitor"
- "Successfully registered default Kubernetes Monitoring UI dashboards."
- "Successfully registered K8Monitor resource metadata"
- "Successfully registered provider with name (keyIndexName) SOMEVALUEHERE and
  id SOMEVALUEHERE"
- "Successfully validated authentication to CEM Event..."
```

If you see `ConfigurationException` errors at start up, you might have provided invalid configuration parameters. Before retrying the deployment, wait for a few minutes while the K8Monitor process reattempts its initialization process in case the Cloud App Management services are not yet ready.

## Kubernetes metrics for thresholds

As soon as you deploy the Kubernetes data collector, incidents are generated for any native Kubernetes events. You can also define your own Kubernetes thresholds that, when breached, open events and generate incidents.

This topic lists the metrics for each Kubernetes resource type that you can use in a threshold definition and provides usage examples.

### Kubernetes Cluster

These are the metrics that are available for use in Kubernetes Cluster thresholds:

- Cluster Name
- CPU: Allocatable Nanocores, Capacity Nanocores, Usage Core Nanoseconds, Usage Millicores
- Deployment Availability Percent
- Ephemeral-Storage: Allocatable Bytes, Capacity Bytes
- File System: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes
- Hugepages-2Mi: Allocatable Bytes, Capacity Bytes
- Memory: Allocatable Bytes, Available Bytes, Capacity Bytes, Major Page Faults, Page Faults, Rss Bytes, Usage Bytes, Usage with Cache Bytes
- Name
- Pods: Allocatable, Capacity, Hosted
- Rlimit: Curproc, Maxpid
- Runtime Image File System: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes
- Stateful Set: Availability, Availability Percent
- Total: Deployments, Deployments Available, Stateful Sets, Stateful Sets Available
- Type

### Kubernetes Container

These metrics are available for use in Kubernetes Container thresholds:

- Cluster Uid
- Container Id
- CPU: Limits Nanocores, Requests Nanocores, Usage Core Nanoseconds
- InitContainer
- Logs: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes
- Memory: Limits Bytes, Major Page Faults, Page Faults, Requests Bytes, Rss Bytes, Usage Bytes, Usage with Cache Bytes
- Name
- Namespace

- Node Id
- Pod Id
- Restart Count
- Rootfs: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes
- Type

**Kubernetes Daemon Set, Kubernetes Deployment, Kubernetes Job, Kubernetes Replication Controller, Kubernetes Replica Set, and Kubernetes Stateful Set**

Use these metrics in thresholds for monitoring the Kubernetes daemon set, deployment, job, replication controller, replica set, or stateful set:

- Annotations
- Available Replicas
- Back off Limit
- Cluster Uid
- Collision Count
- Completion Time
- Completions
- Concurrency Policy
- Creation Timestamp
- Current: Number Scheduled, Replicas, Revision
- Desired Number Scheduled
- Failed (available for Kubernetes Job only)
- Failed Jobs History Limit
- Fully Labeled Replicas
- Generation
- Labels
- Name
- Namespace
- Node Selector (available for Kubernetes Daemon Set only)
- Number: Available, Misscheduled, Ready, Unavailable
- Observed Generation
- Parallelism
- Pod Management Policy
- Ready Replicas
- Replicas
- Revision History Limit
- Schedule
- Service Name
- Start Time
- Succeeded
- Successful Jobs History Limit
- Suspend
- Template Generation
- Update: Revision, Strategy

- Updated: Number Scheduled, Replicas

### **Kubernetes Node**

These metrics are available for use in thresholds for monitoring Kubernetes nodes:

- Allocatable
- Annotations
- Architecture
- Boot Id
- Capacity
- Cluster Uid
- Container Runtime Version
- CPU: Allocatable Nanocores, Capacity Nanocores, Usage Core Nanoseconds, Usage Millicores
- Creation Timestamp
- Ephemeral-Storage: Allocatable Bytes, Capacity Bytes
- External Id
- File System: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes
- Hostname
- Hugepages-2Mi: Allocatable Bytes, Capacity Bytes
- Internal Ip
- Kernel Version
- Kube Proxy Version
- Kubelet: Port, Version
- Labels
- Machine Id
- Memory: Allocatable Bytes, Available Bytes, Capacity Bytes, Major Page Faults, Page Faults, Rss Bytes, Usage Bytes, Usage with Cache Bytes
- Name
- Node Role
- Operating System
- Os Image
- podCIDR
- Pods: Allocatable, Capacity, Hosted
- Rlimit: Curproc, Maxpid
- Runtime Image File System: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes
- System Uuid
- Type
- Unschedulable

### **Kubernetes Pod**

The following metrics are available for use in thresholds for monitoring Kubernetes pods:

- Annotations
- Cluster Uid
- CPU Usage Core Nanoseconds



- Creation Timestamp
- Dns Policy
- Ephemeral-Storage: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes
- Generate Name
- Host: Ip, Network, Pid
- Hostname
- Image Pull Secrets
- Labels
- Memory: Major Page Faults, Page Faults, Rss Bytes, Usage Bytes, Usage with Cache Bytes
- Name
- Namespace
- Network: Received Bytes, Received Errors, Transmitted Bytes, Transmitted Errors
- Node: Id, Name, Selector
- Num Containers
- Phase
- Pod Ip
- Qos Class
- Restart: Count, Policy
- Scheduler Name
- Service: Account, Account Name
- Start Time
- Subdomain
- Termination Grace Period Seconds
- Type

### **Kubernetes Service**

You can use the following metrics to define a Kubernetes Service threshold:

- Annotations
- Browser: Load Time (ms), Type, Version
- Cluster: IP, Uid
- Content Loading Time (ms)
- Creation Timestamp
- Domain Name
- Error Count per Interval
- External Traffic Policy
- Labels
- Latency (ms)
- Load Balancer
- Name
- Namespace
- Page Transfer Time (ms)
- Pod Name
- Ports

- Real User Latency (ms)
- Request: Name, Type
- Resolve Time (ms)
- Selector
- Service Type
- Session Affinity
- Status
- Status Code
- Transaction: Name, Type

Some metrics cannot be used in a threshold definition with multiple AND conditions:

- Request Name and Latency (ms) with Cluster Ip, Creation Timestamp, or Error Count per Interval.
- Labels, Latency (ms), Load Balancer, Name, Namespace, Ports, Request Name

## Uninstalling the Kubernetes data collector

If you no longer want to monitor a resource with the Kubernetes data collector, you can uninstall it. Similarly, if you intend to reinstall a data collector, you must first uninstall it.

### Procedure

Depending on whether you installed the Kubernetes data collector with an Ansible script, with no Ansible script, or without Helm, complete one of these steps on the system where the data collector is installed:

- If you installed the data collector using an Ansible script, enter these kubectl commands:
  - a) Purge the Helm release. If TLS (Transport Layer Security) is not enabled, do not include **--tls** in the command:

```
helm del --purge my_releasename --tls
```

- b) Delete the agent configuration secret:

```
kubectl -n my_namespace delete secret dc-secret
```

- c) Delete the HTTPS secret:

```
kubectl -n my_namespace delete secret ibm-agent-https-secret
```

- d) (Optional) Clean up loaded images:

```
\`
docker rmi -f `docker images | grep icam-k8-monitor | head -1 | awk
'{{print $3}}'`
```

- If you installed the data collector without using an Ansible script, enter these kubectl commands:
  - a) Delete the deployment:

```
kubectl -n my_namespace delete deployment my_releasename-k8monitor
```

- b) Delete the secrets:

```
kubectl -n my_namespace delete secret dc-secret
kubectl -n my_namespace delete secret ibm-agent-https-secret
kubectl -n my_namespace delete secret my_releasename-admintenants
```

- c) (Optional) Clean up loaded images:

```
\`
docker rmi -f `docker images | grep icam-k8-monitor | head -1 | awk
'{{print $3}}'`
```

- If you installed the data collector without Helm, enter these kubectl commands:

```
kubectl delete -f deploy/crds/k8sdc_cr.yaml
kubectl delete -f deploy/operator.yaml
kubectl delete -f deploy/role_binding.yaml
kubectl delete -f deploy/role.yaml
kubectl delete -f deploy/service_account.yaml
kubectl delete -f deploy/crds/k8sdc_crd.yaml
```

Optionally, delete the secrets:

```
kubectl delete secrets dc-secret ibm-agent-https-secret my_pullsecretname
```

## Results

The Kubernetes data collector is uninstalled. Within a few minutes, it no longer shows in the Cloud App Management console.

## Kubernetes data collector issues

Use this topic to review possible causes and solutions to Ansible or Kubernetes data collector installation issues, as well as data retrieval issues.

### Installing klusterlet via Helm

#### Problem

You are installing the klusterlet via helm, and by mistake, deleted the helm and now you can't redeploy the klusterlet since there is already a custom resource definition, for example: 'k8sdcs.ibmcloudapppgmt.com'.

When you try to delete the custom resource definition it hangs.

You get a `Internal service error : rpc error: code = Unknown desc = object is being deleted: customresourcedefinitions.apiextensions.k8s.io "k8sdcs.ibmcloudapppgmt.com" already exists` message.

#### Cause

This error occurs when you attempt to perform the klusterlet install but a custom resource definition already exists and you are unable to delete the custom resource definition.

#### Solution

Check if the custom resource exists by running the command:

```
kubectl get K8sDC -n multicluster-endpoint
```

If it does exist, first patch it using the command:

```
kubectl patch k8sdcs.ibmcloudapppgmt.com -p '{"metadata":{"finalizers":[]}}' --type=merge
Your_CR_name -n multicluster-endpoint
```

You can then delete the custom resource definition, and reinstall the klusterlet via helm.

### Ansible install

#### Problem

You get a `dc-secret already exists` or `ibm-agent-https-secret already exists` message.

#### Cause

This error occurs when you attempt to perform the Ansible install in a namespace where another install was already performed.

#### Solution

If you do not have any running Kubernetes data collector releases in this namespace, delete the existing secret with the following command and run the script again:

```
kubectl -n myNamespace delete secret dc-secret
```

or

```
kubectl -n myNamespace delete secret ibm-agent-https-secret
```

If a running Kubernetes data collector release already exists in this namespace, you need to either remove that release using the procedure in [“Uninstalling the Kubernetes data collector” on page 574](#) and then re-install the data collector, or install this second release into a different namespace.

## Kubernetes data collector install

### Problem

Instead of success initialization indicators in the installation logs (see [“Checking the Kubernetes installation logs” on page 569](#)), you get warning messages.

### Cause

Potential issues that might be the cause:

- Proper ingresses aren't configured on the backend server
- HTTPS is not enabled on the backend server
- Invalid Authentication provided
- Invalid configuration provided, such as the wrong tenantID or ingress (or ingresses). Ensure that you are directing data to the right backend
- Backend services not yet ready
- Backend services struggling
- Unsuccessful collection cycle. This is not critical and could be due to unexpected data or backend services struggling. The data collector reinitializes the cache and tries again after the next interval: After 10 minutes of unsuccessful cycles, the pod will recycle

### Solution

Review the logs for indicators, then review and adjust the settings.

## Dashboard pages show no data

### Problem

No metrics are displayed in the Kubernetes data collector dashboard pages.

### Cause

The local domain cannot be resolved. The K8Monitor component is unable to register dashboards when the IBM Cloud Private cluster doesn't resolve the master node on the DNS (Domain Name System) server.

You can check for an unresolved domain by running the **nslookup** command on the master node (such as **nslookup master-node.cn.ibm.com**). A message that the server can't find the master-node.address confirms that the domain is unresolved.

### Solution

1. Install and configure the NLnet Labs Unbound DNS resolver utility. For more information, see <https://nlnetlabs.nl/documentation/unbound/>.
2. Modify `/etc/resolv.conf` on all cluster VMs: Add the Unbound server as the DNS server.
3. Replace `/etc/resolv.conf` that you modified in step [“2” on page 576](#) on all monitored machines.
4. Restart the cluster's kube-dns pod, which is responsible for the DNS resolution (such as service name and domain name) in the container.
5. Restart the agent. (For more information, see [“Using agent commands” on page 226.](#))

## Runtime data collectors

Runtime data collectors monitor Go, J2SE, Liberty, Node.js, Python, and Ruby applications and are embedded within the application workloads.

### Common topics

Some topics are common to Go, J2SE, Liberty, Node.js, Python, and Ruby data collector.

#### Authorizing the data collector to access Kubernetes resources

To monitor applications that are running in IBM Cloud Private, the service account that you use to configure the runtime data collector must have access to Kubernetes resources through Kubernetes API. Otherwise, you must authorize the service account with appropriate access before you configure the data collector.

#### About this task

The service account that you use to install and configure the data collector must have access to Kubernetes resources. To determine whether the data collector has access to resources, you can use this service account to run the following commands on the Kubernetes master node:

```
kubectl auth can-i list nodes --all-namespaces --as system:serviceaccount:  
namespace:service_account_name  
kubectl auth can-i get nodes --all-namespaces --as system:serviceaccount:  
namespace:service_account_name  
kubectl auth can-i get pods --all-namespaces --as system:serviceaccount:  
namespace:service_account_name  
kubectl auth can-i list services --all-namespaces --as system:serviceaccount:  
namespace:service_account_name  
kubectl auth can-i get services --all-namespaces --as system:serviceaccount:  
namespace:service_account_name  
kubectl auth can-i get configmaps --all-namespaces --as system:serviceaccount:  
namespace:service_account_name  
kubectl auth can-i get deployments --all-namespaces --as system:serviceaccount:  
namespace:service_account_name  
kubectl auth can-i list endpoints --all-namespaces --as system:serviceaccount:  
namespace:service_account_name  
kubectl auth can-i get endpoints --all-namespaces --as system:serviceaccount:  
namespace:service_account_name
```

Where *namespace* is the namespace of your environment and *service\_account\_name* is the name of the service account that you use to configure the data collector. By default, the *service\_account\_name* is default. You must change the *namespace* and *service\_account\_name* with the values that you use.

**Tip:** To determine the existing *service\_account\_name* of a pod, you can run the command **kubectl get po my\_pod\_name -o yaml | grep serviceAccount**, where *my\_pod\_name* is the name of the running pod.

See the following example:

```
kubectl auth can-i list nodes --all-namespaces --as system:serviceaccount:ops-am:default  
kubectl auth can-i get nodes --all-namespaces --as system:serviceaccount:ops-am:default  
kubectl auth can-i get pods --all-namespaces --as system:serviceaccount:ops-am:default  
kubectl auth can-i list services --all-namespaces --as system:serviceaccount:ops-am:default  
kubectl auth can-i get services --all-namespaces --as system:serviceaccount:ops-am:default  
kubectl auth can-i get configmaps --all-namespaces --as system:serviceaccount:ops-am:default  
kubectl auth can-i get deployments --all-namespaces --as system:serviceaccount:ops-am:default  
kubectl auth can-i list endpoints --all-namespaces --as system:serviceaccount:ops-am:default  
kubectl auth can-i get endpoints --all-namespaces --as system:serviceaccount:ops-am:default
```

If you get at least one response of the commands to be no, it means that you do not have required permissions. Do the following steps to grant required service account that is used to set up your application.

#### Procedure

1. Create a ClusterRole yaml file (for example, name it as `lwdc-clusterrole.yaml`) to allow read permission to required Kubernetes resources.

Here is an example:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: lwdc-query
rules:
- apiGroups:
  - ""
  - "apps"
  resources:
  - nodes
  - services
  - configmaps
  - pods
  - deployments
  - endpoints
  verbs:
  - list
  - get
```

2. Run the following command to create the ClusterRole:

```
# kubectl create -f lwdc-clusterrole.yaml
```

3. Create a ClusterRoleBinding yaml file (for example, name it as `lwdc-rolebinding.yaml`) to bind the service account to the ClusterRole that is created in step 1 and 2. This ClusterRole has access permission to query Kubernetes resources in the RBAC mode.

The following example binds the `system:serviceaccount:ops-am:default` account to the specific ClusterRole.

```
kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: lwdc-rolebinding
  namespace: ops-am
subjects:
- kind: User
  name: system:serviceaccount:ops-am:default
  apiGroup: rbac.authorization.k8s.io
roleRef:
  kind: ClusterRole
  name: lwdc-query
  apiGroup: rbac.authorization.k8s.io
```

If you need to grant the access to multiple service accounts in cluster scope, you can also create the ClusterRoleBinding against service account group, for example:

```
kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: lwdc-rolebinding
  namespace: ops-am
subjects:
- kind: Group
  name: system:serviceaccounts
  apiGroup: rbac.authorization.k8s.io
roleRef:
  kind: ClusterRole
  name: lwdc-query
  apiGroup: rbac.authorization.k8s.io
```

4. Run the following command:

```
# kubectl create -f clusterrolebinding.yaml
```

### Obtaining the server configuration information

During data collector deployment, the server information must be provided so that you can configure the data collector to connect to the appropriate server. You must configure the connection from the data

collector to the Cloud App Management server by downloading a configuration package from the Cloud App Management console

### About this task

After the Cloud App Management server is deployed, the server information is provided as a package for download from the Cloud App Management console.

### Procedure

1. Download the data collector configuration package:
  - a) Log in to the Cloud App Management console, click the **Get Started** link on the Welcome page, then select **Administration > Integrations**.
  - b) Click the **New integration** button.
  - c) In the **Standard monitoring agents** section, select the **ICAM Data Collectors Configure** button.
  - d) Select **Download file** and specify the directory where you want to save the compressed data collector configuration package, `ibm-cloud-apm-dc-configpack.tar`.
2. Extract the `ibm-cloud-apm-dc-configpack.tar` file to get the `global.environment` file and the keyfiles.

**Note:** These files contain all the variables and values that are required by data collectors to connect to the server.

### What to do next

Configure the server connection for data collectors and update the application deployment. See [“Monitoring Liberty applications in Kubernetes environment” on page 591](#), [“Monitoring Node.js applications in Kubernetes environment” on page 601](#), [“Monitoring J2SE applications in Kubernetes environment” on page 585](#), and [“Monitoring Python applications in Kubernetes environment” on page 608](#).

## Configuring Go application monitoring

The Go data collector can provide you with visibility and control of your Go applications, and help you ensure optimal performance and efficient use of resources. You can reduce and prevent application crashes and slowdowns around the clock, as the data collector assists you in detecting, diagnosing and isolating performance issues.

The Go data collector helps you to manage the performance and availability of the following:

- Go applications in Kubernetes environments
- Local Go applications

For the detailed system requirements of Go data collector, see [system requirements for Go data collector](#).

### Downloading the Go data collector

You can download the Go data collector package from Passport Advantage.

### About this task

To download the Go data collector package, complete the following steps:

### Procedure

1. Review the part numbers and components to download. For more information, see [Part numbers](#).
2. Extract the package to get the Go data collector (`go_datacollector.tgz`) package file by running the following command:

```
tar xzf appMgtDataCollectors_2019.4.0.tar.gz
cd appMgtDataCollectors_2019.4.0
```

```
tar xzf app_mgmt_runtime_dc_2019.4.0.tar.gz
cd app_mgmt_runtime_dc_2019.4.0
```

## Monitoring Go applications in Kubernetes environment

Before you monitor Go applications in Kubernetes environment, you must connect the data collector to the server by creating a secret. Then, you update your application deployment to monitor the Go applications.

### Before you begin

- Check whether your service account has access to Kubernetes resources. For more information, see [“Authorizing the data collector to access Kubernetes resources”](#) on page 577.
- Ensure that you downloaded the configuration package to obtain the server information. For more information, see [“Obtaining the server configuration information”](#) on page 578.
- Check whether you downloaded the Go data collector package from Passport Advantage. For more information, see [“Downloading the Go data collector”](#) on page 579.

### About this task

To enable the Go data collector, you need to update the application code to import the Go data collector package, rebuild the application with the vendor files of the Go data collector, and create the secret based on the configuration package.

### Procedure

1. Update the Go application by importing the Go data collector module in the application main file.

```
import (
    _ "github.ibm.com/APM/godc"
)
```

2. Rebuild the Go application with vendor files of the Go data collector.

- a) Extract the Go data collector package file `go_datacollector.tgz`:

```
tar xzf go_datacollector.tgz
```

- b) Merge the vendor files with the vendor files of the application.
- c) Build the application with new vendor files.

```
go build -mod=vendor
```

- d) Create the Kubernetes secret based on the configuration package. Go to the `ibm-cloud-apm-dc-configpack` directory where you extract the configuration package in [“Obtaining the server configuration information”](#) on page 578, and run the following command to create a secret to connect to the server, for example, name it as `icam-server-secret`.

```
kubectl -n default create secret generic icam-server-secret \
--from-file=ibm-cloud-apm-dc-configpack/keyfiles/keyfile.jks \
--from-file=ibm-cloud-apm-dc-configpack/keyfiles/keyfile.p12 \
--from-file=ibm-cloud-apm-dc-configpack/keyfiles/keyfile.kdb \
--from-file=ibm-cloud-apm-dc-configpack/keyfiles/ca.pem \
--from-file=ibm-cloud-apm-dc-configpack/keyfiles/cert.pem \
--from-file=ibm-cloud-apm-dc-configpack/keyfiles/key.pem \
--from-file=ibm-cloud-apm-dc-configpack/global.environment
```

3. Update the application yaml file to mount the secret. See the following example.

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: MyGoApp
  labels:
    app: MyGoApp
spec:
  selector:
```



```

matchLabels:
  app: MyGoApp
  pod: MyGoApp
replicas: 1
template:
  metadata:
    name: MyGoApp
    labels:
      app: MyGoApp
      pod: MyGoApp
  spec:
    containers:
      - name: MyGoApp
        image: mycluster.icp:8500/default/MyGoApp:v1
        imagePullPolicy: Always
        ports:
          - containerPort: 3000
            protocol: TCP
        env:
          - name: APPLICATION_NAME
            value: "MyGoApp"
        volumeMounts:
          - name: global-environment
            mountPath: /opt/ibm/apm/serverconfig
    volumes:
      - name: global-environment
        secret:
          secretName: icam-server-secret
          optional: true

```

4. Redeploy the Go application.

### Monitoring on-premises Go applications

You can configure the Go data collector to monitor the on-premises Go applications and then send monitoring data to the Cloud App Management server.

#### Before you begin

- Ensure that you downloaded the configuration package to obtain the server information. For more information, see [“Obtaining the server configuration information” on page 578](#).
- Check whether you downloaded the Go data collector package. For more information, see [“Downloading the Go data collector” on page 579](#).

#### About this task

To enable the Go data collector, you need to update the application code to import the Go data collector package, rebuild the application with the vendor files of the Go data collector, and configure the data collector.

#### Procedure

1. Update the Go application by importing the Go data collector module in the application main file.

```

import (
    _ "github.ibm.com/APM/godc"
)

```

2. Rebuild the Go application with vendor files of the Go data collector.

a) Extract the Go data collector package file `go_datacollector.tgz`:

```
tar xzf go_datacollector.tgz
```

b) Merge the vendor files with the vendor files of the application.

c) Build the application with new vendor files.

```
go build -mod=vendor
```

3. Configure the Go data collector. Put the file `global.environment` and the folder `keyfiles` into the root of Go application.

4. Redeploy the Go application.

### Customizing the Go data collector

You can set the variables to change the default behavior of the Go data collector.

#### User-defined environment variables for the Go data collector

For Go monitoring in Kubernetes environments, set the following variables in the application yml file.

For on-premises applications, set the following variables as environment variables or in the `config.properties` file.

Table 79. Supported user-defined environment variables for Go monitoring		
Variable name	Value	Description
OPENTRACING_ENABLED	False	By default, the OpenTracing function is enabled. You can disable OpenTracing by setting the environment variable to false.  Example: <pre>- name: OPENTRACING_ENABLED   value: false</pre>
OpenTracing sampling: • JAEGER_SAMPLER_TYPE • JAEGER_SAMPLER_PARAM	The default sampler type is probabilistic, and the default sampler param is 0.01, which means that 1 in 100 traces will be sampled. You can set it to other values.  The Go data collector supports three sampler types: <ul style="list-style-type: none"><li>• Constant (JAEGER_SAMPLER_TYPE=const)</li><li>• Probabilistic (JAEGER_SAMPLER_TYPE=probabilistic)</li><li>• Rate Limiting (JAEGER_SAMPLER_TYPE=ratelimiting)</li></ul> For more information, see <a href="#">Sampling</a> .	When the OpenTracing function is enabled, you can set the OpenTracing sampler type and param. Example: <pre>- name: JAEGER_SAMPLER_TYPE   value: probabilistic - name: JAEGER_SAMPLER_PARAM   value: "0.1"</pre>

Table 79. Supported user-defined environment variables for Go monitoring (continued)

Variable name	Value	Description
LATENCY_SAMPLER_PARAM	Any value between 0 and 1	<p>The default value is 0.1, which means getting one request out of 10 requests. The value must be between 0 and 1. The value of 0 means no latency data will be collected. The value of 1 means no sampler and all requests data will be collected.</p> <p>Example:</p> <pre data-bbox="935 495 1463 583">- name: LATENCY_SAMPLER_PARAM   value: "0.2"</pre>
GODC_LOG_LEVEL	<ul style="list-style-type: none"> <li>• PANIC</li> <li>• FATAL</li> <li>• ERROR</li> <li>• WARN</li> <li>• INFO</li> <li>• DEBUG</li> <li>• TRACE</li> </ul>	<p>The default value is INFO. You can set it to the following levels:</p> <ul style="list-style-type: none"> <li>• PANIC: PanicLevel level, the highest level of severity. Logs and then calls panic with the message passed to Debug, Info, and so on.</li> <li>• FATAL: FatalLevel level. Logs and then calls <code>logger.Exit(1)</code>. It will exit even if the logging level is set to Panic.</li> <li>• ERROR: ErrorLevel level. Logs. Used for errors that should definitely be noted. Commonly used for hooks to send errors to an error tracking service.</li> <li>• WARN: WarnLevel level. Non-critical entries that deserve eyes.</li> <li>• DEBUG: Only useful debug information is printed in the log, for example, collected data, data that are sent to server, and server response.</li> <li>• INFO: InfoLevel level. General operational entries about what's going on inside the application.</li> <li>• DEBUG: DebugLevel level. Usually only enabled when debugging. Very verbose logging.</li> <li>• TRACE: TraceLevel level. Designates finer-grained informational events than the Debug.</li> </ul> <p>Example:</p> <pre data-bbox="935 1640 1463 1696">export GODC_LOG_LEVEL=level</pre>

Table 79. Supported user-defined environment variables for Go monitoring (continued)

Variable name	Value	Description
GODC_LOG_TO_CONSOLE	true	Set GODC_LOG_TO_CONSOLE to dump the log to console, otherwise, the output will be dumped to log files under the root of application. The log files look like godc.log and restclient.log.  Example: <pre>export GODC_LOG_TO_CONSOLE=true</pre>

### Uninstalling the Go data collector from your application

To uninstall the Go data collector, roll back the changes that you have made to your application and then update the application deployment.

#### Procedure

- To uninstall the Go data collector in Kubernetes environment:
  - In the application main file, remove `_ "github.ibm.com/APM/godc"` from `import`.
  - Remove configpack reference by removing the secret reference and the corresponding mount volume in application yaml file:

```
volumeMounts:
  - mountPath: /opt/ibm/apm/serverconfig
    name: serverconfig
```

and

```
volumes:
  - name: global-environment
    secret:
      secretName: icam-server-secret
    optional: true
```

- Rebuild the Docker image and redeploy your application.
- To uninstall the Go data collector in on-premises environment:
    - In the application main file, remove `_ "github.ibm.com/APM/godc"` from `import`.
    - Rebuild the application and redeploy it.

## Configuring J2SE application monitoring

You can use J2SE data collector to monitor the cloud-based Java applications. The J2SE data collector is a greenfield runtime data collector, which is available with the IBM Cloud App Management Advanced package.

The J2SE data collector helps you to manage the performance and availability of Spring Boot based applications and other stand-alone Java applications in Kubernetes environments.

You can configure the J2SE data collector to send data to the Cloud App Management server.

For the detailed system requirements of J2SE data collector, see [system requirements for J2SE data collector](#).

### Downloading the J2SE data collector

You can download the J2SE data collector package for the target system from Passport Advantage.

#### About this task

To download the J2SE data collector from Passport Advantage, follow these steps:

## Procedure

1. Review the part numbers and components to download. For more information, see: [Part numbers](#).
2. Unpack the greenfield package to get the latest J2SE data collector (`j2se_datacollector.tgz`) by running the following command:

```
tar -xzf appMgtDataCollectors_2019.4.0.tar.gz
cd appMgtDataCollectors_2019.4.0
tar -xzf app_mgmt_runtime_dc_2019.4.0.tar.gz
cd app_mgmt_runtime_dc_2019.4.0
```

## Monitoring J2SE applications in Kubernetes environment

Before you monitor J2SE applications in IBM Cloud Private or OpenShift, you must connect the data collector to the IBM Cloud App Management server by creating a secret. Then, you can update your application deployment to monitor the J2SE applications.

### Before you begin

- Check whether your service account has access to Kubernetes resources. For more information, see [“Authorizing the data collector to access Kubernetes resources”](#) on page 577.
- Check whether the J2SE data collector is available for download from Passport Advantage. For more information, see [“Downloading the J2SE data collector”](#) on page 584.
- Check whether you downloaded the configuration package to obtain the server information. For more information, see [“Obtaining the server configuration information”](#) on page 578.

### About this task

Configure the J2SE data collector to the server by creating a secret. Then, update the application deployment to use the Docker file that you build. You can create a secret by using the `global.environment` file and `keyfiles` that are extracted from the Cloud App Management configuration package. Then, you can mount this secret when you deploy the application as a Kubernetes deployment.

## Procedure

1. Go to the `ibm-cloud-apm-dc-configpack` directory where you extract the configuration package in [“Obtaining the server configuration information”](#) on page 578, and run the following command to create a secret to connect to the server, for example, name it as `icam-server-secret`.

```
kubectl -n my_namespace create secret generic icam-server-secret \
--from-file=keyfiles/keyfile.jks \
--from-file=keyfiles/keyfile.p12 \
--from-file=keyfiles/keyfile.kdb \
--from-file=global.environment
```

Where `my_namespace` is the namespace where you want to create the secret. If you want to create the secret in the default namespace, remove `-n my_namespace` from the command.

2. Update the Docker file of your J2SE application to include J2SE data collector details. Following is the sample of a Docker file:

```
WORKDIR dc_dir
COPY /j2se_datacollector.tgz dc_dir
RUN tar -xzf dc_dir/j2se_datacollector.tgz
COPY /silent_config_j2se_dc.txt dc_dir/bin/
COPY app.jar /opt/
RUN dc_dir/bin/config_dc.sh -silent
RUN chmod +x dc_dir/runtime/j2se${applicationAlias}.${localhostName}.${applicationAlias}/
dcstartup.sh
RUN mkdir /opt/logs
ENTRYPOINT [ "sh", "-c", "dc_dir/runtime/j2se${applicationAlias}.${localhostName}.${
{applicationAlias}/dcstartup.sh" ]
```

Where:

- `dc_dir` is the directory where the J2SE data collector installer is located, for example, `/opt/j2se_dc`.
- `j2se_datacollector.tgz` is the file name of the downloaded J2SE data collector installation package.
- `app.jar` is the J2SE application that is packaged in a runnable JAR.
- `silent_config_j2se_dc.txt` is the name of the silent configuration file for configuring the J2SE data collector.
- The `j2se_datacollector.tgz`, `app.jar`, and `silent_config_j2se_dc.txt` file must be in the same directory as your Docker file. This is required because Docker needs all files in the context of the Docker build. For more information about writing the Docker file, see: [Dockerfile reference](#).
- You can get the `applicationAlias` value from the silent configuration file.

Following is the example of a silent configuration file for J2SE data collector:

```

JAVA_HOME=/opt/ibm/java/jre
# APPLICATION_TYPE, 1:Java Application, 2:Jetty Server
APPLICATION_TYPE=1
APPLICATION_HOME=/opt/EmployeeWeb-1.0-SNAPSHOT.jar
MAIN_CLASS=com.venk.springboot.EmployeeWebApp
APPLICATION_ALIAS=j2seApp
TT_STATUS=TRUE
DD_STATUS=FALSE
MT_STATUS=TRUE

```

3. Build and tag the new Docker image of the application and push this new image to the private registry. Also, ensure that you include the docker registry and the docker group when you build and push the image, as shown here:

```

docker build -t <docker_registry>/<docker_group>/<application_image_name>:<image_tag>
docker push <docker_registry>/<docker_group>/<application_image_name>:<image_tag>

```

For example,

```

docker build -t mycluster.icp:8500/default/my_app_image:latest
docker push mycluster.icp:8500/default/my_app_image:latest

```

4. Open your application deployment `yaml` file to use the new Docker image and update the `volumeMounts` and `Volumes` section by adding the following:

```

volumeMounts:
- name: global-environment
  mountPath: /opt/ibm/apm/serverconfig
volumes:
- name: global-environment
  secret:
    secretName: icam-server-secret
    optional: true

```

Where:

- `/opt/ibm/apm/serverconfig` is the fixed value to store the files in the docker container.
  - `icam-server-secret` is the name of the secret that is created in step 1.
5. If you are working with a local application deployment `yaml` file, then you must run the following command for the changes to take effect:

```

kubectl create -f application_deployment_yaml_file -n my_namespace

```

## Monitoring on-premises J2SE applications

You can configure the J2SE data collector to monitor the on-premises J2SE applications running on stand-alone Docker containers, VMs, or physical nodes and then send monitoring data to the Cloud App Management server.

### Before you begin

- Check that you downloaded the configuration package to obtain the server information, for more information, see [“Obtaining the server configuration information” on page 578](#).
- Check whether you downloaded the J2SE data collector package. For more information, see [“Downloading the J2SE data collector” on page 584](#).

### Procedure

1. Extract the `j2se_datacollector.tgz` file that you get from [“Downloading the J2SE data collector” on page 584](#).

```
tar -xvf path-of-j2se_datacollector.tgz
```

2. With `config_dc.sh` script under J2SE data collector directory, and the `ibm-cloud-apm-dc-configpack.tar` that is downloaded in [“Obtaining the server configuration information” on page 578](#), run the following command to apply server configuration to the monitored application:

```
path-of-j2se_datacollector.tgz/bin/config_dc.sh [-silent <silent_file>]
```

**Note:** If `silent_file` is not provided, the default `silent_config_j2se_dc.txt` under the same directory of `config_dc.sh` will be used.

3. Copy `configpack` files into J2SE data collector directory:

```
cp ./ibm-cloud-apm-dc-configpack/global.environment /opt/j2se_dc/itcamdc/etc
cp ./ibm-cloud-apm-dc-configpack/keyfiles/keyfile.jks /opt/j2se_dc/itcamdc/etc
```

4. Run `mkdir ../logs`.

5. Run the following command:

```
chmod +x path-of-j2se_datacollector.tgz/runtime/
j2seAPPLICATION_ALIAS.HOSTNAME.APPLICATION_ALIAS/dcstartup.sh
```

Where:

- `APPLICATION_ALIAS` is the value that is configured in `APPLICATION_ALIAS` item in `silent_file`.
- The default value of `HOSTNAME` is `localhost`.

6. Run the following command:

```
sh -c path-of-j2se_datacollector.tgz/runtime/
j2seAPPLICATION_ALIAS.HOSTNAME.APPLICATION_ALIAS/dcstartup.sh
```

Where:

- `APPLICATION_ALIAS` is the value that is configured in `APPLICATION_ALIAS` item in `silent_file`.
- The default value of `HOSTNAME` is `localhost`.

**Note:** If you run the J2SE application in a docker container, you can start the application in background by running the following command:

```
docker run -d -p <docker_port:app_port> j2seapp:1.0
```

Where `docker_port` is the docker port exposed, `app_port` is the application port, for example, **10090:8080**.

7. If the J2SE application runs in a docker container, rebuild your docker container with J2SE data collector and `configpack` installed by running .

```
docker build -t < application image name >:< image tag >
```

Docker file example:

```
FROM ibmjava:8
WORKDIR /opt/j2se_dc
COPY /j2se_datacollector.tgz /opt/j2se_dc
RUN tar -xvf /opt/j2se_dc/j2se_datacollector.tgz
COPY /silent_config_j2se_dc.txt /opt/j2se_dc/bin/
COPY /EmployeeWeb-1.0-SNAPSHOT.jar /opt/
COPY /ibm-cloud-apm-dc-configpack/global.environment /opt/j2se_dc/itcamdc/etc/
COPY /ibm-cloud-apm-dc-configpack/keyfiles/keyfile.jks /opt/j2se_dc/itcamdc/etc/
RUN /opt/j2se_dc/bin/config_dc.sh -silent
COPY /dcstartup.sh /opt/j2se_dc/runtime/j2sej2seApp.localhost.j2seApp/
COPY /del_slf4.sh /opt/j2se_dc/runtime/j2sej2seApp.localhost.j2seApp/
RUN chmod +x /opt/j2se_dc/runtime/j2sej2seApp.localhost.j2seApp/dcstartup.sh
RUN mkdir /opt/logs
ENTRYPOINT [ "sh", "-c", "/opt/j2se_dc/runtime/j2sej2seApp.localhost.j2seApp/dcstartup.sh" ]
```

Where:

- /opt/j2se\_dc is *path-of-j2se\_datacollector.tgz*.
- EmployeeWeb-1.0-SNAPSHOT.jar is the application name.
- j2seApp is APPLICATION\_ALIAS item value in *silent\_file*.

### Customizing the J2SE data collector

You can set the variables to change the default behavior of the J2SE data collector.

### User-defined environment variables for the J2SE data collector

For J2SE monitoring in IBM Cloud Private or OpenShift, set the following variables in the application yaml file.

For on-premises applications, set the following variables as environment variables or in the config.properties file.

Variable name	Value	Description
OPENTRACING_ENABLED	"False"	By default, the J2SE data collector enables OpenTracing function. You can disable OpenTracing by setting the environment variable to false.  Example: <pre>- name: OPENTRACING_ENABLED   value: "false"</pre>
OpenTracing sampling: • JAEGER_SAMPLER_TYPE • JAEGER_SAMPLER_PARAM	The default sampler type is "probabilistic", and the default sampler param is 0.01, which means that 1 in 100 traces will be sampled. You can set it to other values. For more information, see <a href="#">Sampling</a> .	When the OpenTracing function is enabled, you can set the OpenTracing sampler type and param. Example: <pre>- name: JAEGER_SAMPLER_TYPE   value: "probabilistic" - name: JAEGER_SAMPLER_PARAM   value: "0.1"</pre>



Table 80. Supported user-defined environment variables for J2SE monitoring (continued)

Variable name	Value	Description
LATENCY_SAMPLER_PARAM	Any value between 0 and 1	<p>The default value is 0.1, which means getting 1 request out of 10 requests. The value must be between 0 and 1. The value of 0 means no latency data will be collected. The value of 1 means no sampler and all requests data will be collected.</p> <p>Example:</p> <pre>- name: LATENCY_SAMPLER_PARAM   value: "0.2"</pre>

### Uninstalling the J2SE data collector from your application

You can uninstall the J2SE data collector in Kubernetes environment or on-premises environment.

#### Procedure

- If you deploy your J2SE data collector in Kubernetes environments, do the following steps to uninstall:

a) Roll back changes to the Docker file by doing the following steps:

a. Remove the following lines from the Docker file:

```
WORKDIR /opt/j2se_dc
RUN tar -xf /opt/j2se_dc/j2se_datacollector.tgz
COPY /silent_config_j2se_dc.txt /opt/j2se_dc/bin/
COPY /EmployeeWeb-1.0-SNAPSHOT.jar /opt/
RUN /opt/j2se_dc/bin/config_dc.sh -silent
RUN chmod +x /opt/j2se_dc/runtime/j2sej2seApp.localhost.j2seApp/dcstartup.sh
RUN mkdir /opt/logs
ENTRYPOINT [ "sh", "-c", "/opt/j2se_dc/runtime/j2sej2seApp.localhost.j2seApp/
dcstartup.sh" ]
```

b. Add the following lines to the Docker file:

```
WORKDIR /opt
COPY /appstartup.sh /opt/
ENTRYPOINT [ "sh", "-c", "/opt/appstartup.sh" ]
```

Example of the Docker file after rolling back:

```
FROM ibmjava:8
WORKDIR /opt
COPY /EmployeeWeb-1.0-SNAPSHOT.jar /opt/
COPY /appstartup.sh /opt/
ENTRYPOINT [ "sh", "-c", "/opt/appstartup.sh" ]
```

**Note:** You need to manually create the `appstartup.sh` file which contains only one line. See the following example:

```
/opt/ibm/java/jre/bin/java" -jar /opt/EmployeeWeb-1.0-SNAPSHOT.jar
```

b) Rebuild the Docker images.

c) Remove the secret reference and the corresponding mount volume in the application yml file and redeploy your application.

```
volumeMounts:
- mountPath: /opt/ibm/apm/serverconfig
  name: serverconfigCopy
volumes:
- name: global-environment
  secret:
```

```
secretName: icam-server-secret
optional: true
```

- d) Roll back changes to the application yaml file and redeploy your application.
- If you deploy your J2SE data collector in on-premises docker environment, do the following steps to uninstall:
  - a) Roll back changes to the Docker file by doing the following steps:
    - a. Remove the following lines from the Docker file:

```
WORKDIR /opt/j2se_dc
RUN tar -xf /opt/j2se_dc/j2se_datacollector.tgz
COPY /silent_config_j2se_dc.txt /opt/j2se_dc/bin/
COPY /EmployeeWeb-1.0-SNAPSHOT.jar /opt/
RUN /opt/j2se_dc/bin/config_dc.sh -silent
RUN chmod +x /opt/j2se_dc/runtime/j2sej2seApp.localhost.j2seApp/dcstartup.sh
RUN mkdir /opt/logs
ENTRYPOINT [ "sh", "-c", "/opt/j2se_dc/runtime/j2sej2seApp.localhost.j2seApp/
dcstartup.sh" ]
```

- b. Add the following lines to the Docker file:

```
WORKDIR /opt
COPY /appstartup.sh /opt/
ENTRYPOINT [ "sh", "-c", "/opt/appstartup.sh" ]
```

Example of the Docker file after rolling back:

```
FROM ibmjava:8
WORKDIR /opt
COPY /EmployeeWeb-1.0-SNAPSHOT.jar /opt/
COPY /appstartup.sh /opt/
ENTRYPOINT [ "sh", "-c", "/opt/appstartup.sh" ]
```

**Note:** You need to manually create the `appstartup.sh` file which contains only one line. See the following example:

```
/opt/ibm/java/jre/bin/java" -jar /opt/EmployeeWeb-1.0-SNAPSHOT.jar
```

- b) Rebuild the Docker images and redeploy the application.
  - If you deploy your J2SE data collector on bare-metal node or VMs, disable J2SE data collector, stop or kill server process.

## Configuring Liberty application monitoring

You can use the Liberty data collector to monitor your Liberty applications. The Liberty data collector is a runtime data collector that runs within the Liberty profile.

The Liberty data collector helps you to manage the performance and availability of Java-based microservices or Liberty applications in Kubernetes environments.

You can configure the data collector to send data to the Cloud App Management server.

For the detailed system requirements of Liberty data collector, see [system requirements for Liberty data collector](#).

### (Conditional) Downloading the Liberty data collector

The Liberty data collector is available in the [WebSphere Liberty Repository](#) and it can be automatically downloaded if your local system can access this online public repository. If your firewall rules do not allow connection to the WebSphere Liberty Repository, you can download it from another system that has access. Alternatively, you can download the data collector from Passport Advantage.

#### Procedure

To download the Liberty data collector, complete the download from one of the following locations:

- WebSphere Liberty Repository

If your environment doesn't have access to the WebSphere Liberty repository where the Liberty application container is built, complete the following steps:

- a. Download the extension pack as a .esa file from another system that can access [Liberty data collector](#) in the WebSphere Liberty Repository.
- b. Copy the downloaded .esa file to a temporary directory on your local system where the Liberty application is running.

If your environment has access to the [WebSphere Liberty repository](#) where the Liberty application container is built, complete the following step:

- a. Add the following to the Docker file, so that the .esa file is downloaded automatically:

```
RUN /opt/ibm/wlp/bin/installUtility install ibmAppMetricsForJava-1-2-1 --acceptLicense
```

- **Passport Advantage**

- a. To download from Passport Advantage, review the part numbers and components to download, for more information, see: [Part numbers](#).
- b. Unpack the greenfield package to get the latest liberty data collector (javametrics.liberty.icam-1.2.1.esa).
- c. Enter the following:

```
tar xzf appMgtDataCollectors_2019.4.0.tar.gz
cd appMgtDataCollectors_2019.4.0
tar xzf app_mgmt_runtime_dc_2019.4.0.tar.gz
cd app_mgmt_runtime_dc_2019.4.0
```

## Monitoring Liberty applications in Kubernetes environment

Before you monitor Liberty applications in IBM Cloud Private or OpenShift, you must connect the data collector to the server by creating a secret. Then, you update your application deployment to monitor the Liberty applications.

### Before you begin

If your service account doesn't have access to Kubernetes resources, see: [“Authorizing the data collector to access Kubernetes resources”](#) on page 577.

The Liberty data collector is available to be automatically downloaded from [WebSphere Liberty Repository](#) during configuration. If your firewall rules do not allow connection to this open repository, download the Liberty data collector from another system that has access. However, if you don't want to download the data collector from the public repository due to company policy and you would prefer to download from Passport Advantage, for more information, see [“\(Conditional\) Downloading the Liberty data collector”](#) on page 590.

Check that you downloaded the configuration package to obtain the server information, for more information, see [“Obtaining the server configuration information”](#) on page 578.

### About this task

Configure the data collector to the server by creating the secret. Then, update the application deployment to use the Docker file that you build.

### Procedure

1. Go to the `ibm-cloud-apm-dc-configpack` directory where you extract the configuration package in [“Obtaining the server configuration information”](#) on page 578, and run the following command to create a secret to connect to the server, for example, name it as `icam-server-secret`.

```
kubectl -n my_namespace create secret generic icam-server-secret \
--from-file=keyfiles/keyfile.jks \
--from-file=keyfiles/keyfile.p12 \
```

```
--from-file=keyfiles/keyfile.kdb \  
--from-file=global.environment
```

Where *my\_namespace* is the namespace where you want to create the secret. If you want to create the secret in the default namespace, remove `-n my_namespace` from the command.

2. Create a `silent_config_liberty_dc.txt` silent configuration file in the same directory as your Dockerfile and add the following lines:

```
JAVA_HOME=/opt/ibm/java/jre  
LIBERTY_HOME=path_to_liberty_home  
SERVER_NAME=*  
ADD_XMX_SIZE=True  
SERVER_TYPE=ICAM  
CONFIGPACK_PATH=/opt/ibm-cloud-apm-dc-configpack.tar
```

Where:

- *JAVA\_HOME* is the Java home that is used by liberty applications. The default value is `/opt/ibm/java/jre`.
- *LIBERTY\_HOME* is the directory where the liberty application is installed. For WebSphere Liberty, the default value is `/opt/ibm/wlp`. For Open Liberty, the default value is `/opt/ol/wlp`.
- *SERVER\_NAME* is the name of the liberty servers that are monitored by the data collector. You can separate the server names with a space character. The `*` character shows all the servers installed are monitored. The default value is `*`.
- *ADD\_XMX\_SIZE* allows you to allocate an extra 512 M memory for all the monitored servers. The value is **True** or **False**. The default value is **True**.
- *SERVER\_TYPE* is the type of monitoring server to which the data collector is connected. The value is **ICAM** or **APM**. The default value is **ICAM**.
- *CONFIGPACK\_PATH* is the absolute path of the configuration package.

**Note:**

- The Cloud App Management server and the Cloud APM server provide different ways to get the configuration package file. When the liberty application is running in IBM Cloud Private and the data collector is connecting to the Cloud App Management server, this variable is ignored.

3. Update the Dockerfile of your Liberty application. You must have the write access to the server folder.

- If you use Open Liberty, add the following lines to your Dockerfile:

```
COPY path_to_esa_file /opt/  
RUN mkdir -p /opt/ol/wlp/usr/extension/lib/features/  
RUN cd /tmp && unzip /opt/javametrics.liberty.icam-1.2.1.esa  
  && mv /tmp/wlp/liberty_dc /opt/ol/wlp/usr/extension/ && mv /tmp/OSGI-INF/SUBSYSTEM.MF  
  /opt/ol/wlp/usr/extension/lib/features/javametrics.liberty.icam-1.2.1.mf  
COPY path_to_silent_file /opt/ol/wlp/usr/extension/liberty_dc/bin/  
RUN /opt/ol/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent
```

- If you use WebSphere Liberty and your environment has access to the WebSphere Liberty repository where the Liberty application container is built, add the following commands to your Dockerfile:

```
RUN chmod 777 liberty_server_dir  
RUN /opt/ibm/wlp/bin/installUtility install  
  ibmAppMetricsForJava-1.2.1  
  --acceptLicense  
RUN /opt/ibm/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent
```

The **chmod 777** command grants you the write access to the Liberty server directory, for example, `/opt/ibm/wlp/usr/server/defaultServer`. The **installUtility** command enables you to download the data collector automatically from the WebSphere Liberty repository.

- If you use WebSphere Liberty and your environment doesn't have access to the WebSphere Liberty repository where the Liberty application container is built, add the following commands to your Dockerfile:

```
RUN chmod 777 liberty_server_dir
COPY path_to_esa_file /opt/
RUN /opt/ibm/wlp/bin/installUtility install --acceptLicense /opt/javametrics.liberty.icam-1.2.1.esa
RUN /opt/ibm/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent
```

You must download the esa file because it is not downloaded automatically by using the **installUtility** command.

Where:

- *path\_to\_esa\_file* is the relative path of the downloaded `javametrics.liberty.icam-1.2.1.esa` file to the current directory. The `javametrics.liberty.icam-1.2.1.esa` file must be in same directory as your Docker file or in a sub-directory of your Docker file location. This is required because Docker needs all files in the context of the Docker build. For more information about writing the Docker file, see: [Dockerfile reference](#).
- **installUtility** is the liberty tool in the bin directory of the Liberty home directory. It is used to install the Liberty extension pack.
- `/opt/ibm/wlp/` is the default path to the WebSphere Liberty home directory. It can be changed accordingly.
- `config_unified_dc.sh` is the configuration script that is used to configure the data collector. It runs in silent mode with the **-silent** parameter, and it reads the `liberty_dc_home/bin/silent_config_liberty_dc.txt` default config file.

**Note:** In the `liberty_dc_home/bin` directory, you can see a default `silent_config_liberty_dc.txt` configuration file, where **LIBERTY\_HOME** is set to `/opt/ibm/wlp`. This setting is right for WebSphere Liberty, so the configuration script can work well. But for Open Liberty, the **LIBERTY\_HOME** parameter should be set to `/opt/ol/wlp`. You need to use the configuration file that you created in step “2” on page 592 and add the **COPY path\_to\_silent\_file liberty\_dc\_home/bin/** command to your Dockerfile. In this way, the configuration script can find the customized configuration file and run with it.

4. Build and tag the new Docker image of the application and push this new image to the private registry.

Ensure that you include the docker registry and the docker group when you build and push the image, as shown here:

```
docker build -t <docker_registry>/<docker_group>
/<application_image_name>:<image_tag> .
docker push <docker_registry>/<docker_group>/<application_image_name>:<image_tag>
```

Example:

```
docker build -t mycluster.icp:8500/default/my_app_image:latest .
docker push mycluster.icp:8500/default/my_app_image:latest
```

5. Open your application deployment yml file to use the new Docker image and update the `volumeMounts` and `Volumes` section by adding the following lines:

```
volumeMounts:
- name: global-environment
  mountPath: /opt/ibm/apm/serverconfig
volumes:
- name: global-environment
  secret:
    secretName: icam-server-secret
  optional: true
```

Where:

- `/opt/ibm/apm/serverconfig` is the fixed value to store the files in the docker container.

- `icam-server-secret` is the name of the secret that is created in step “1” on page 591.

If you are working with a local application deployment `yaml`, you must run the following command for your changes to take effect:

```
kubectl create -f application_deployment_yaml_file
-n my_namespace
```

## Monitoring Microclimate-based Liberty applications in IBM Cloud Private

By default, applications that are created by using Microclimate include the **Microclimate > App Monitor** in-built monitoring dashboards. If these applications are deployed in IBM Cloud Private, in production, you should upgrade to the Cloud App Management server to get polyglot application monitoring, and alerting and analytics across your hybrid cloud applications. If you have access to the Cloud App Management server, you can configure the Microclimate-based Liberty applications easily to be monitored by IBM Cloud App Management.

### Before you begin

Ensure that you have installed and configured the following:

- Microclimate
- Cloud App Management server

### About this task

You must delete the old Cloud APM data collector configuration from the `Dockerfile` that is generated in Microclimate before you configure the Microclimate-based Liberty applications to be monitored by the server.

### Procedure

Deleting the old Cloud APM data collector configuration

1. In your application directory, remove the following lines from the `Dockerfile`:

```
RUN installUtility install --acceptLicense defaultServer
&& installUtility install --acceptLicense
  apmDataCollector-7.4
  ENV
  LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/lib:/opt/ibm/wlp/usr/
  extension/liberty_dc/toolkit/lib/1x8266 \
  JVM_ARGS="$JVM_ARGS -agentlib:am_ibm_16=defaultServer -
  Xbootclasspath/p:/opt/ibm/wlp/usr/extension/liberty_dc/
  toolkit/lib/bcm-bootstrap.jar -
  Xverbosegclog:/logs/gc.log,1,10000 -verbosegc -
  Djava.security.policy=/opt/ibm/wlp/usr/extension/liberty_dc/
  itcamdc/etc/datacollector.policy -
  Dliberty.home=/opt/ibm/wlp"
```

Configuring the Liberty data collector

2. Enable your application with the Cloud App Management data collector.

- If you use **Open Liberty**, add the following lines to your `Dockerfile`:

```
COPY path_to_esf_file /opt/
RUN mkdir -p /opt/ol/wlp/usr/extension/lib/features/
RUN cd /tmp && unzip /opt/javametrics.liberty.icam-1.2.1.esf
  && mv /tmp/wlp/liberty_dc /opt/ol/wlp/usr/extension/ && mv /tmp/OSGI-INF/SUBSYSTEM.MF
  /opt/ol/wlp/usr/extension/lib/features/javametrics.liberty.icam-1.2.1.mf
COPY path_to_silent_file /opt/ol/wlp/usr/extension/liberty_dc/bin/
RUN /opt/ol/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent
```

- If you use **WebSphere Liberty**, add the following lines to the `Dockerfile`:

```
RUN chmod 777 liberty_server_dir
RUN /opt/ibm/wlp/bin/installUtility install
```

```
ibmAppMetricsForJava-1.2.1
--acceptLicense
RUN /opt/ibm/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent
```

This enables your application container to include the Liberty data collector. If the Cloud App Management configuration isn't provided, the data collector remains in a disabled state.

3. To enable the Liberty data collector to communicate with the Cloud App Management server at a later stage, follow step “1” on page 591 and step “5” on page 593 in: [Monitoring Liberty applications in IBM Cloud Private](#).

### Monitoring on-premises Liberty applications

You can configure the Liberty data collector to monitor the on-premises Liberty applications running on stand-alone Docker containers, VMs or physical nodes and then send monitoring data to the Cloud App Management server.

#### Before you begin

- Check that you downloaded the configuration package to obtain the server information, for more information, see [“Obtaining the server configuration information”](#) on page 578.
- Check whether you downloaded the Liberty data collector package. For more information, see [“\(Conditional\) Downloading the Liberty data collector”](#) on page 590.

#### About this task

Configure the Liberty data collector by using the `config_unified_dc.sh` file under the Liberty data collector directory, for example, `liberty_home/usr/extension/liberty_dc/bin`, and the `ibm-cloud-apm-dc-configpack.tar` file that you downloaded from the server.

#### Procedure

- To configure the Liberty data collector for VMs or physical nodes, do the following steps:
  - a) Go to the `liberty_home/bin` directory and run the following command to install the data collector:

```
./installUtility install --acceptLicense path_to_dc_package
```

where `path_to_dc_package` is the full path to the `.esa` file that you downloaded. For example, `/opt/javametrics.liberty.icam-1.2.1.esa`.

**Note:** The `installUtility` command is not applicable to Open Liberty. If you use Open Liberty, copy the `javametrics.liberty.icam-1.2.1.esa` file to the `/opt` directory and run the following commands:

```
mkdir -p /opt/ol/wlp/usr/extension/lib/features/
cd /tmp
unzip /opt/javametrics.liberty.icam-1.2.1.esa
mv /tmp/wlp/liberty_dc /opt/ol/wlp/usr/extension/
mv /tmp/OSGI-INF/SUBSYSTEM.MF /opt/ol/wlp/usr/extension/lib/features/
javametrics.liberty.icam-1.2.1.mf
```

- b) Run the following command to apply server configuration to the monitored application:

```
./config_unified_dc.sh [-silent <silent_file>]
```

#### Note:

- If `<silent_file>` is not provided, the `silent_config_liberty_dc.txt` file under the same directory of `config_unified_dc.sh` will be used.
- The path of `ibm-cloud-apm-dc-configpack.tar` is provided in `silent_file`, for example, `CONFIGPACK_PATH=/opt/ibm-cloud-apm-dc-configpack.tar`.

- c) Restart your Liberty application.

- To configure the Liberty data collector for standalone Docker environment, do the following steps:

- a) Create a `silent_config_liberty_dc.txt` silent configuration file in the same directory as your Dockerfile and add the following lines:

```
JAVA_HOME=/opt/ibm/java/jre
LIBERTY_HOME=/opt/ibm/wlp
SERVER_NAME=*
ADD_XMX_SIZE=True
SERVER_TYPE=ICAM
CONFIGPACK_PATH=/opt/ibm-cloud-apm-dc-configpack.tar
```

Where:

- `JAVA_HOME` is the Java home that is used by liberty applications. The default value is `/opt/ibm/java/jre`.
- `LIBERTY_HOME` is the directory where the liberty application is installed. The default value is `/opt/ibm/wlp`.
- `SERVER_NAME` is the name of the liberty servers that are monitored by the data collector. You can separate the server names with a space character. The `*` character shows all the servers installed are monitored. The default value is `*`.
- `ADD_XMX_SIZE` allows you to allocate an extra 512 M memory for all the monitored servers. The value is **True** or **False**. The default value is **True**.
- `SERVER_TYPE` is the type of monitoring server to which the data collector is connected. The value is **ICAM** or **APM**. The default value is **ICAM**.
- `CONFIGPACK_PATH` is the absolute path of the configuration package.

- b) Update the specific variable values according to the Liberty server settings.

- c) If you use Open Liberty, add the following lines to the Dockerfile of your Liberty application:

```
COPY path_to_esa_file /opt/
RUN mkdir -p /opt/ol/wlp/usr/extension/lib/features/
RUN cd /tmp && unzip /opt/javametrics.liberty.icam-1.2.1.esa
&& mv /tmp/wlp/liberty_dc /opt/ol/wlp/usr/extension/
&& mv /tmp/OSGI-INF/SUBSYSTEM.MF /opt/ol/wlp/usr/extension/lib/features/
javametrics.liberty.icam-1.2.1.mf
COPY path_to_silent_file /opt/ol/wlp/usr/extension/liberty_dc/bin/
RUN /opt/ol/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent
```

Where:

- `path_to_esa_file` is the relative path for the downloaded `javametrics.liberty.icam-1.2.1.esa` file to the current directory. For example, `/tmp/javametrics.liberty.icam-1.2.1.esa`.
- `path_to_silent_file` is the relative path for the `silent_config_liberty_dc.txt` file to the current directory. For example, `/tmp/silent_config_liberty_dc.txt`.

- d) If you use WebSphere Liberty, add the following lines to the Dockerfile of your Liberty application:

```
COPY path_to_esa_file /opt/
RUN /opt/ibm/wlp/bin/installUtility install --acceptLicense /opt/
javametrics.liberty.icam-1.2.1.esa
COPY path_to_silent_file /opt/ibm/wlp/usr/extension/liberty_dc/bin/
RUN /opt/ibm/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent
```

- e) Build the new docker image.

```
docker build -t <application image name>:<image tag> .
```

- f) Start your Liberty application with the new docker image.



## Customizing the Liberty data collector

You can set the variables to change the default behavior of the Liberty data collector.

### User-defined environment variables for the Liberty data collector

For Liberty monitoring in IBM Cloud Private or OpenShift, set the following variables in the application yaml file.

For on-premises applications, set the following variables as environment variables or in the `config.properties` file.

Variable name	Value	Description
OPENTRACING_ENABLED	"False"	By default, the Liberty data collector enables OpenTracing function. You can disable OpenTracing by setting the environment variable to false.  Example: <pre>- name: OPENTRACING_ENABLED   value: "false"</pre>
OpenTracing sampling: <ul style="list-style-type: none"><li>• JAEGER_SAMPLER_TYPE</li><li>• JAEGER_SAMPLER_PARAM</li></ul>	The default sampler type is "probabilistic", and the default sampler param is 0.01, which means that 1 in 100 traces will be sampled. You can set it to other values. For more information, see <a href="#">Sampling</a> .	When the OpenTracing function is enabled, you can set the OpenTracing sampler type and param. Example: <pre>- name: JAEGER_SAMPLER_TYPE   value: "probabilistic" - name: JAEGER_SAMPLER_PARAM   value: "0.1"</pre>
LATENCY_SAMPLER_PARAM	Any value between 0 and 1	The default value is 0.1, which means getting 1 request out of 10 requests. The value must be between 0 and 1. The value of 0 means no latency data will be collected. The value of 1 means no sampler and all requests data will be collected.  Example: <pre>- name: LATENCY_SAMPLER_PARAM   value: "0.2"</pre>

### DEM for Liberty applications

After you configure the Liberty data collector, you can enable DEM to collect data about how actual users interact with and experience web applications.

#### Before you begin

Prerequisites:

- DEM supports monitoring kube services that are exposed by Ingress only. Make sure that the following settings are enabled in your Kubernetes environment.

- If SSL Passthrough is enabled in ingress, ensure that X-FORWARDED-FOR is enabled. Set below annotation to ingress to pass client IP to X-FORWARDED-FOR header.

```
ingress.kubernetes.io/configuration-snippet: |
    proxy_set_header X-Forwarded-For $proxy_protocol_addr;
```

- If you are using ingress rewriting rules, ensure X-Original-URI header is supported to pass original URI. By default it is enabled.
- Make sure the “Kubernetes data collector” on page 557 is installed and enabled. Otherwise, the DEM service cannot be found in IBM Cloud App Management portal. For more information, see “Kubernetes data collector” on page 557.
- Make sure the Liberty data collector that you deploy is up-to-date. For more information, see “Monitoring Liberty applications in Kubernetes environment” on page 591.
- OpenTracing monitoring must be enabled. By default it is enabled in Liberty monitoring.

### About this task

To enable DEM after you configure the Liberty data collector, do the following steps:

### Procedure

1. Open your application deployment yaml file, and add the following environment variables.

```
containers:
  - env:
    - name: IBM_APM_RUM_ENABLED
      value: "true"
```

**Note:** If you want to disable DEM, set the value to false.

2. Apply your application deployment yaml file.
3. Optional: If you have configured OpenTracing sampling settings of the Liberty data collector, DEM reads the sampler type and param. If you want to change the sampling settings, open the application deployment yaml file and modify the lines in the **env:** section, for example,

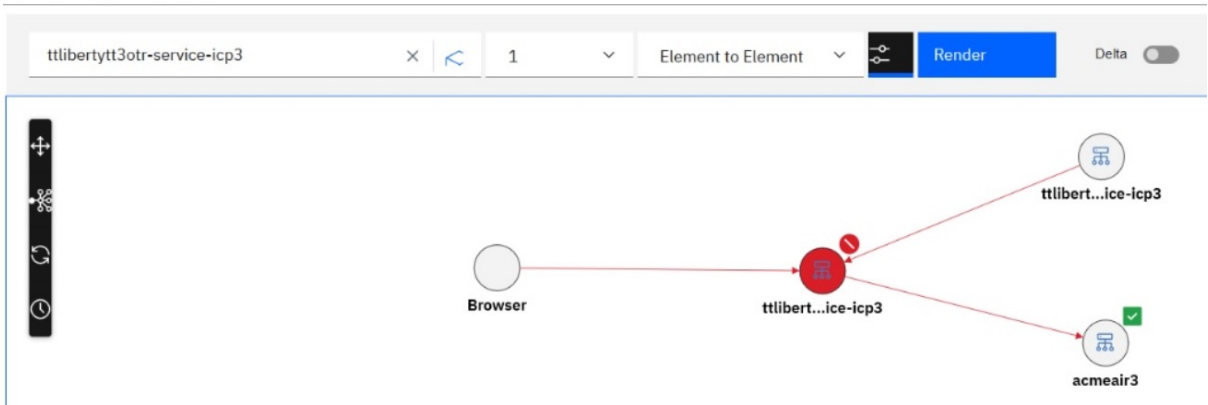
```
- name: JAEGER_SAMPLER_TYPE
  value: "rateLimiting"
- name: JAEGER_SAMPLER_PARAM
  value: "10"
```

For more information, see “Customizing the Liberty data collector” on page 597.

### What to do next

Launch a web request, and then do the following steps to verify whether DEM is successfully enabled:

1. In the Cloud App Management console, click the **Resources** tab.
2. Find **Kubernetes Service** from the **All resource types** list and click to open it.
3. Browse the Resource list, click the resource name that you have enabled DEM, and open the resource dashboard.
4. Use one of the following methods to verify:
  - Check whether **browser** is displayed in **Service dependencies** section.



- Check whether Browser is listed in the **Related resources** widget.

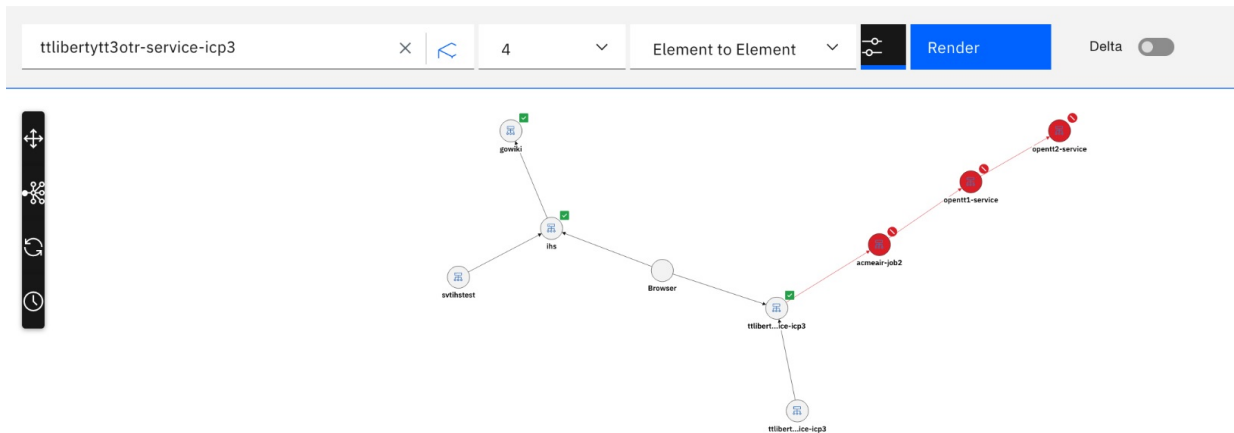
**Related resources**

Filter

Status	Relation	Resource	Type
✓ Clear	communicatesW	<a href="#">rumdcdev-864768f74b-n67pn (devCluster:default)</a>	Kubernetes Pod
✓ Clear	communicatesW	<a href="#">rumdcdev (devCluster:default)</a>	Kubernetes Ingress
◆ Indeter...	dependsOn/in	<a href="#">Browser</a>	Browser

5. Click to drill down browser to check whether you can get detailed browser data.

**Known limitation:** If multiple services are connected to the Browser in topology, no browser data can be displayed.



### Uninstalling the Liberty data collector from your application

You can uninstall the Liberty data collector in Kubernetes environment or on-premises environment.

#### Procedure

- If you deploy your Liberty data collector in Kubernetes environments, do the following steps to uninstall the Liberty data collector:

- a) Remove configpack reference by removing the secret reference and the corresponding mount volume in application yaml file:

```
volumeMounts:
  - mountPath: /opt/ibm/apm/serverconfig
    name: serverconfig
```

and

```
volumes:
  - name: global-environment
    secret:
      secretName: icam-server-secret
    optional: true
```

- b) If you use Open Liberty, remove the following lines from the Dockerfile to roll back changes:

```
COPY path_to_esa_file /opt/
RUN mkdir -p /opt/ol/wlp/usr/extension/lib/features/
RUN cd /tmp && unzip /opt/javametrics.liberty.icam-1.2.1.esa && mv /tmp/wlp/
liberty_dc /opt/ol/wlp/usr/extension/ && mv /tmp/OSGI-INF/SUBSYSTEM.MF /opt/ol/wlp/usr/
extension/lib/features/javametrics.liberty.icam-1.2.1.mf
COPY path_to_silent_file /opt/ol/wlp/usr/extension/liberty_dc/bin/
RUN /opt/ol/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent
```

- c) If you use WebSphere Liberty and your environment doesn't have access to the WebSphere Liberty repository where the Liberty application container is built, remove the following lines from the Dockerfile to roll back changes:

```
RUN chmod 777 liberty_server_dir
COPY path_to_esa_file /opt/
RUN /opt/ibm/wlp/bin/installUtility install --acceptLicense /opt/javametrics.liberty.icam-
1.2.1.esa
RUN /opt/ibm/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent
```

- d) If you use WebSphere Liberty and your environment has access to the WebSphere Liberty repository where the Liberty application container is built, remove the following lines from the Dockerfile to roll back changes:

```
RUN chmod 777 liberty_server_dir
RUN /opt/ibm/wlp/bin/installUtility install
ibmAppMetricsForJava-1.2.1
--acceptLicense
RUN /opt/ibm/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent
```

- e) Rebuild the Docker images and redeploy your application.

- If you deploy your Liberty data collector in on-premises docker environment, do the following steps to uninstall the Liberty data collector:

- a) If you use Open Liberty, remove the following lines from the Dockerfile to roll back changes:

```
COPY path_to_esa_file /opt/
RUN mkdir -p /opt/ol/wlp/usr/extension/lib/features/
RUN cd /tmp && unzip /opt/javametrics.liberty.icam-1.2.1.esa
&& mv /tmp/wlp/liberty_dc /opt/ol/wlp/usr/extension/ && mv /tmp/OSGI-INF/SUBSYSTEM.MF
/opt/ol/wlp/usr/extension/lib/features/javametrics.liberty.icam-1.2.1.mf
COPY path_to_silent_file /opt/ol/wlp/usr/extension/liberty_dc/bin/
RUN /opt/ol/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent
```

- b) If you use WebSphere Liberty, remove the following lines from the Dockerfile to roll back changes:

```
COPY path_to_esa_file /opt/
RUN /opt/ibm/wlp/bin/installUtility install --acceptLicense /opt/
javametrics.liberty.icam-1.2.1.esa
COPY path_to_silent_file /opt/ibm/wlp/usr/extension/liberty_dc/bin/
RUN /opt/ibm/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent
```

- c) Rebuild the Docker images and redeploy the application.

- If you deploy your Liberty data collector on bare-metal node or VMs, stop Liberty server and run the following command to uninstall the Liberty data collector:

```
./unconfig_liberty_dc.sh [-silent <silent_file>]
```

## Configuring Node.js application monitoring

You can use the Node.js data collector to monitor your Node.js-based applications. The Node.js data collector provides you with visibility and control of your Node.js applications, and helps you to ensure optimal performance and efficient use of your resources. You can reduce and prevent application crashes and slowdowns around the clock, as the data collector assists you in detecting, diagnosing, and isolating performance issues.

The Node.js data collector helps you to manage the performance and availability of Node.js and Microclimate-based applications in Kubernetes environments.

For the detailed system requirements of Node.js data collector, see [system requirements for Node.js data collector](#).

### Installing the Node.js data collector

Depending on your environment and whether you can access the internet, you can install the Node.js data collector by using different procedures.

#### Procedure

To install the data collector, complete one of the following procedures:

- If your environment can access the internet:
  - To update the `package.json`, add `"appmetrics": "^4.0.0"` as a dependency.
  - To update the main Node application file, add the `require('appmetrics');` to the beginning of the file.
- If your environment can't access the internet or if your company policy doesn't allow you to download packages from open source, complete the following steps:
  - Unpack the data collectors package according to your Node.js Runtime version, for example, for `appMgtDataCollectors_2019.4.0.tar.gz`:

```
tar xzf appMgtDataCollectors_2019.4.0.tar.gz
cd appMgtDataCollectors_2019.4.0
tar xzf app_mgmt_runtime_dc_2019.4.0.tar.gz
cd app_mgmt_runtime_dc_2019.4.0
tar xzf nodejs_datacollector_2019.4.0.tgz
tar xzf ibmapm-greenfield-v8-1x64.tgz
```

For more information, see [“Obtaining the server configuration information”](#) on page 578.

- Copy or move the `ibmapm` folder that is created in step 1 to the root folder of your Node application. The root folder is the folder that contains the Node application file.

```
mv ibmapm application_root_folder/ibmapm
```

- Add `require('./ibmapm');` to the first line of your application entry file.

### Monitoring Node.js applications in Kubernetes environment

Before you monitor Node.js applications in IBM Cloud Private or OpenShift, you must connect the data collector to the server by creating a secret. Then you update your application deployment to monitor the Node.js applications.

#### Before you begin

If your service account doesn't have access to Kubernetes resources, see: [“Authorizing the data collector to access Kubernetes resources”](#) on page 577.

Check that you downloaded the configuration package to obtain the server information, for more information, see [“Obtaining the server configuration information”](#) on page 578.

Ensure that you installed the data collector, for more information, see [Installing the Node.js data collector](#).

### About this task

You can create a secret for the `global.environment` file and the keyfiles that are extracted from the Cloud App Management configuration package. Then, you mount this secret when you deploy the application as a Kubernetes deployment.

### Procedure

1. Go to the `ibm-cloud-apm-dc-configpack` directory where you extract the configuration package in [“Obtaining the server configuration information”](#) on page 578, and run the following command to create a secret to connect to the server, for example, name it as `icam-server-secret`.

```
kubectl -n my_namespace create secret generic icam-server-secret \
--from-file=keyfiles/keyfile.p12 \
--from-file=global.environment
```

Where *my\_namespace* is the namespace where you want to create the secret. If you want to create the secret in the default namespace, remove `-n my_namespace` from the command.

2. Update the Docker file of your Node.js application to get the write access to the work directory by adding the following line:

```
RUN chmod 777 nodejs_dir
```

Where *nodejs\_dir* is the home directory of your Node.js application, for example, `/var/apps/acmeair-nodejs`.

3. Build and tag the new docker image of the application and push the new image to the private registry. For example, in the directory where the Docker is located, run the following command:

```
docker build -t <application image name>:<image tag>
```

4. To update the application `yaml` file to mount the secret, complete the following steps:

- a. Add the volume mount information to the `Containers:` object in the application deployment `yaml` file as shown here:

```
volumeMounts:
  - mountPath: /opt/ibm/apm/serverconfig
    name: serverconfig
```

- b. Add the volume information to the `Spec:` object in the application deployment `yaml` file as shown here:

```
volumes:
  - name: global-environment
    secret:
      secretName: icam-server-secret
      optional: true
```

Example of a `Yaml` file that is updated:

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: acmeair
  labels:
    app: acmeair
spec:
  selector:
    matchLabels:
      app: acmeair
      pod: acmeair
```

```

replicas: 1
template:
  metadata:
    name: acmeair
    labels:
      app: acmeair
      pod: acmeair
  spec:
    containers:
    - name: acmeair
      image: mycluster.icp:8500/default/acmeair:v1
      imagePullPolicy: Always
      ports:
      - containerPort: 3000
      protocol: TCP
      env:
      - name: KNJ_LOG_TO_FILE
        value: "true"
      - name: KNJ_LOG_LEVEL
        value: "debug"
      - name: APPLICATION_NAME
        value: "acmeair"
      volumeMounts:
      - name: serverconfig
        mountPath: /opt/ibm/apm/serverconfig
    volumes:
    - name: global-environment
      secret:
        secretName: icam-server-secret
        optional: true

```

5. Update the application yaml file to use the new docker image.

### Monitoring Microclimate-based Node.js applications in IBM Cloud Private

If you have Microclimate-based Node.js applications that you want to monitor in IBM Cloud Private, you must first set up a connection between the data collector and the Cloud App Management server. Then you update your application deployment to monitor the Microclimate-based Node.js applications.

#### Before you begin

If your service account doesn't have access to Kubernetes resources, see: [“Authorizing the data collector to access Kubernetes resources” on page 577](#).

Check that you downloaded the configuration package to obtain the server information, for more information, see [“Obtaining the server configuration information” on page 578](#).

You must check the `appmetrics` version before you install the data collector. The Node.js application that is created by Microclimate requires `appmetrics` automatically. Check the `appmetrics` version and if the version is 4.0.0 or later, then the Node.js data collector is included already. If the version is less than 4.0.0, you must upgrade the `appmetrics` to version 4.0.0 or later. For more information, see [Installing the Node.js data collector](#).

#### Procedure

Follow the procedure to monitor Node.js applications in IBM Cloud Private here: [“Monitoring Node.js applications in Kubernetes environment” on page 601](#). You can create a secret to configure the Cloud App Management server by using a `global.environment` file and `keyfiles` that are extracted from the Cloud App Management configuration package. Then, you mount this secret when you deploy the application as a Kubernetes deployment.

## Monitoring on-premises Node.js applications

You can configure the Node.js data collector to monitor the on-premises Node.js applications running on stand-alone Docker containers, VMs, or physical nodes (xLinux only). The Node.js data collector sends monitoring data to the Cloud App Management server.

### Before you begin

Check that you downloaded the configuration package to obtain the server information, for more information, see [“Obtaining the server configuration information” on page 578](#).

Ensure that you installed the data collector, for more information, see [Installing the Node.js data collector](#).

### About this task

Configure the Node.js data collector using the `apply_configpack.sh` file that you extract from `nodejs_datacollector_2019.4.0.tgz`, and `ibm-cloud-apm-dc-configpack.tar` file that you download from the server.

### Procedure

- To configure the Node.js data collector for VMs or physical nodes, run the following command to apply server configuration to the monitored application:

```
./apply_configpack.sh path_configpack [application folder]
```

Where *application folder* is the directory where you run `node <yourapp.js>`, default value is current folder.

- To configure the Node.js data collector for stand-alone Docker environment, do the following steps:
  - a) Go to the directory where you extract the configuration package in [“Obtaining the server configuration information” on page 578](#), and run the following command:

```
./apply_configpack.sh path_configpack [application folder]
```

Where *application folder* is the directory where you run `node <yourapp.js>`, default value is current folder.

- b) Update the Docker file of your Node.js application by adding the following line to get the write access to the work directory:

```
RUN chmod 777 nodejs_dir
```

Where *nodejs\_dir* is the home directory of your Node.js application, for example, `/var/apps/acmeair-nodejs`.

- c) Rebuild your docker container with Node.js dc (ibmapm) and configpack installed.

```
docker build -t <application image name>:<image tag>
```

### Customizing the Node.js data collector

You can set the variables to change the default behavior of the Node.js data collector.

### User-defined environment variables for the Node.js data collector

For Node.js monitoring in IBM Cloud Private or OpenShift, set the following variables in the application yml file.

For on-premises applications, set the following variables as environment variables or in the `config.properties` file.



Table 82. Supported user-defined environment variables for Node.js monitoring

Variable name	Value	Description
OPENTRACING_ENABLED	"False"	<p>By default, the Node.js data collector enables OpenTracing function. You can disable OpenTracing by setting the environment variable to false.</p> <p>Example:</p> <pre>- name: OPENTRACING_ENABLED   value: "false"</pre> <p><b>Note:</b> If you install Node.js data collector by requiring "appmetrics", the OpenTracing is disabled by default.</p>
OPENTRACING_SAMPLER	The default value is 0.01, which means that 1 in 100 traces will be sampled. You can set it to other values.	<p>When the OpenTracing function is enabled, you can set the OpenTracing sampler rate.</p> <p>Example:</p> <pre>- name: OPENTRACING_SAMPLER   value: "0.1"</pre>
LATENCY_SAMPLER_PARAM	Any value between 0 and 1	<p>The default value is 0.1, which means getting 1 request out of 10 requests. The value must be between 0 and 1. The value of 0 means no latency data will be collected. The value of 1 means no sampler and all requests data will be collected.</p> <p>Example:</p> <pre>- name: LATENCY_SAMPLER_PARAM   value: "0.2"</pre>

### Uninstalling the Node.js data collector

To uninstall the Node.js data collector, roll back the changes that you have made to your application and then update the application deployment.

### Procedure

- If the Node.js application in IBM Cloud Private is deployed by using Microclimate, complete the following steps:
  - a) Complete one of the following steps:
    - To uninstall the Node.js data collector that was set up without internet access, edit the main file of your Node.js application to remove the following line:
 

```
require('./ibmapm');
```
    - To uninstall the Node.js data collector that was set up with internet access, edit the main file of your Node.js application to remove the following line:
 

```
require('appmetrics');
```
  - b) Push your project to a new repository that the Microclimate pipeline is monitoring.
- If the Node.js application is not deployed by using Microclimate, complete the following steps:
  - a) Complete one of the following steps:

- To uninstall the Node.js data collector that was set up without internet access, edit the main file of your Node.js application to remove the following line:

```
require('./ibmapm');
```

- To uninstall the Node.js data collector that was set up with internet access, edit the main file of your Node.js application to remove the following line:

```
require('appmetrics');
```

#### b) Remove configpack reference.

- For Kubernetes environment, remove the secret reference and the corresponding mount volume in the application yaml file:

```
volumeMounts:  
  - mountPath: /opt/ibm/apm/serverconfig  
    name: serverconfig
```

```
volumes:  
  - name: global-environment  
    secret:  
      secretName: icam-server-secret  
      optional: true
```

- For local on-premise or docker container environment, run the following command:

```
rm -f global.environment  
rm -f keyfile.p12
```

#### c) Remove all Node.js data collector resources from the application.

- To uninstall the Node.js data collector that was set up without internet access, edit the main file of your Node.js application to remove the following line:

```
rm -rf ibmapm
```

- To uninstall the Node.js data collector that was set up with internet access, edit the main file of your Node.js application to remove the following dependencies line in package.json of your application:

```
"appmetrics": "^5.0.0"
```

#### d) Apply the changes to make the uninstalling take effect.

- In local on-premise environment, delete node\_modules folder from the home directory of your application, and then run the **npm install** command to install the application dependencies.
- In Docker container environment (whether Kubernetes or not), you need to rebuild your docker image.

## Configuring Python application monitoring

You can use the Python data collector to monitor your Python applications. Through detecting, diagnosing, and isolating performance issues, the Python data collector helps you ensure optimal performance and efficient use of resources, reduce, and prevent application crashes and slowdowns around the clock.

You can configure the Python data collector to connect to the Cloud App Management server. The Python data collector helps you to manage the performance and availability of the following:

- Python applications with Django or Flask frameworks in Kubernetes environments
- Local Python applications with Django or Flask frameworks

**Important:** The Python data collector can monitor the applications running on the internal web servers of Django or Flask frameworks, mod\_wsgi/apache httpd and uWSGI.

## System requirements

For the detailed system requirements of Python data collector, see [system requirements for Python data collector](#).

## Prerequisites

Make sure that you install `psutil` to your Python application. Otherwise, you might get an error when deploying the Python data collector:

```
gcc -pthread -Wno-unused-result -Wsign-compare -DNDEBUG -O2 -g -pipe -Wall -Wp,-
D_FORTIFY_SOURCE=2 -fexceptions -fstack-protector-strong --param=ssp-buffer-size=4 -grecord-
gcc-switches -m64 -mtune=generic -D_GNU_SOURCE -fPIC -fwrapv -fPIC -DPSUTIL_VERSION=430 -
I/usr/include/python3.6m -c psutil/_psutil_linux.c -o build/temp.linux-x86_64-3.6/psutil/
_psutil_linux.o
psutil/_psutil_linux.c:12:20: fatal error: Python.h: No such file or directory
#include <Python.h>
^
compilation terminated.
error: command 'gcc' failed with exit status 1
```

Depending on your Python version (2.7 or 3) and Linux operation system type, run the proper tool to install `psutil`.

- If it is Ubuntu or Debian, and Python version is 2, run the following command:

```
sudo apt-get install python-dev
```

- If it is Ubuntu or Debian, and Python version is 3, run the following command:

```
sudo apt-get install python3-dev
```

- If it is CentOS or RHEL, and Python version is 2.7, run the following command:

```
sudo yum install python-devel
```

- If it is CentOS or RHEL, and Python version is 3, run the following command:

```
sudo yum install python3-devel
```

## Downloading the Python data collector

You can download the Python data collector package from Passport Advantage.

## About this task

To download the Python data collector package, complete the following steps:

## Procedure

1. Review the part numbers and components to download, for more information, see: [Part numbers](#).
2. Extract the package to get the latest Python data collector (`ibm_python_datacollector.tgz`) by running the following command:

```
tar xzf appMgtDataCollectors_2019.4.0.tar.gz
cd appMgtDataCollectors_2019.4.0
tar xzf app_mgmt_runtime_dc_2019.4.0.tar.gz
cd app_mgmt_runtime_dc_2019.4.0
```

## Monitoring Python applications in Kubernetes environment

Before you monitor Python applications in IBM Cloud Private or OpenShift, you must connect the data collector to the server by creating a secret. Then, you update your application deployment to monitor the Python applications.

### Before you begin

- Check whether your service account has access to Kubernetes resources. For more information, see [“Authorizing the data collector to access Kubernetes resources” on page 577](#).
- Ensure that you downloaded the configuration package to obtain the server information. For more information, see [“Obtaining the server configuration information” on page 578](#).
- Check whether you downloaded the Python data collector package from Passport Advantage and placed it with Python application Dockerfile in the same directory. For more information, see [“Downloading the Python data collector” on page 607](#).

### About this task

To configure the Python data collector, pass the IBM Cloud App Management server configuration through secret. You can create a secret for the `global.environment` file and keyfiles that are extracted from the IBM Cloud App Management configuration package. Then, you can mount this secret when you deploy the application as a Kubernetes deployment.

### Procedure

1. Update the Dockerfile to add the following lines to install Python data collector for your Python application, and get write access to the root directory.

```
ADD ibm_python_datacollector.tgz root_of_application
RUN chmod 777 root_of_application
RUN pip install --no-index --find-links=root_of_application/python_dc ibm_python_dc
```

Where `root_of_application` is the Python application root directory of the context of the build (the Dockerfile).

2. From IBM Cloud App Management V2019.4.0, you can use the Python data collector to monitor Python applications running on the web server uWSGI V1.9.0 or later versions. But some configurations are needed to make the Python data collector work well.
  - If you start uWSGI without threads, the threads that are generated by your application and the Python data collector will never run and thus the Python data collector cannot work normally. You must enable uWSGI threads by adding the `--enable-threads` option in the `uwsgi` command. This option is applied automatically when you specify the `--threads` option to configure the number of threads.
  - If you run the `uwsgi` command with the `--master` option, Python scripts and modules are preloaded in the parent master process, and worker processes are forked from the parent master process. In addition, background threads that are created in the master process are killed in worker processes. To make the Python data collector work normally, you must add the `--lazy-apps` option in the `uwsgi` command to use the lazy loading mode.

Command example:

```
uwsgi --enable-threads --master --lazy-apps --processes 4 --http :8002 --wsgi-
file=flask_hello.py --callable app
```

**Note:** You can also add the options in a `.ini` configuration file and run the `uwsgi` command with the `.ini` file. For more information, see uWSGI documents.

3. Integrate the installed data collector in your Python application:

- If your application is based on Django V1.10 or later versions, open `settings.py` of your Django application, and add the following content into first line of section **MIDDLEWARE** in that file:

```
'ibm_python_dc.kpg_dc_django.ResourceMiddleware',
```

- If your application is based on Django V1.9 or older versions, add the following content into the first line of the section **MIDDLEWARE\_CLASS** in that file:

```
'ibm_python_dc.kpg_dc_django.ResourceMiddleware',
```

- If your application is based on Flask, add the Python data collector wsgi middleware in your Python application file, for example, if you run **export FLASK\_APP=run.py**, then edit the `run.py` file and ensure that the Python data collector wsgi middleware are added in front of other middlewares. Example:

```
from flask import Flask
from flask_restful import Api
from api.board import Article
from api.auth import Login, Register, RefreshToken
from middleware import Test
from werkzeug.middleware.dispatcher import DispatcherMiddleware

api.add_resource(Login, '/login')
api.add_resource(Register, '/register')

from ibm_python_dc.kpg_dc_wsgi import ResourceMiddleware
app.wsgi_app = ResourceMiddleware(app.wsgi_app)

app.wsgi_app = DispatcherMiddleware(serve_frontend, {
    '/test': test,
    '/admin': admin,
})

.....
.....
```

4. Go to the `ibm-cloud-apm-dc-configpack` directory where you extract the configuration package in [“Obtaining the server configuration information”](#) on page 578, and run the following command to create a secret to connect to the server, for example, name it as `icam-server-secret`.

```
kubectl -n my_namespace create secret generic icam-server-secret \
--from-file=keyfiles/keyfile.jks \
--from-file=keyfiles/keyfile.p12 \
--from-file=keyfiles/keyfile.kdb \
--from-file=global.environment
```

Where `my_namespace` is the namespace where you want to create the secret. If you want to create the secret in the default namespace, remove `-n my_namespace` from the command.

5. Update the application yml file to mount the secret. See the following example.

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: djangoapp
  labels:
    app: djangoapp
spec:
  selector:
    matchLabels:
      app: djangoapp
      pod: djangoapp
  replicas: 1
  template:
    metadata:
      name: djangoapp
      labels:
        app: djangoapp
        pod: djangoapp
    spec:
      containers:
        - name: djangoapp
          image: mycluster.icp:8500/default/djangoapp:v1
```

```

imagePullPolicy: Always
ports:
- containerPort: 8001
protocol: TCP
env:
- name: NAMESPACE_DEFAULT
  value: "default"
- name: KPG_LOG_TOCONSOLE
  value: "True"
- name: KPG_LOG_LEVEL
  value: "INFO"
volumeMounts:
- name: serverconfig
  mountPath: /opt/ibm/apm/serverconfig
volumes:
- name: global-environment
  secret:
    secretName: icam-server-secret
    optional: true

```

6. Build the new Docker image.

7. Update the application yaml file to use the new Docker image.

### Monitoring on-premises Python applications

The Python data collector supports monitoring Python applications running on stand-alone Docker containers, VMs, or physical nodes (xLinux only).

#### About this task

#### Procedure

- To monitor on-premises Python applications on VMs or physical nodes, follow the instructions in [“Monitoring on-premises Python applications running on VMs or physical nodes”](#) on page 610.
- To monitor on-premises Python applications on stand-alone Docker containers, follow the instructions in [“Monitoring on-premises Python applications running in individual Docker container”](#) on page 612.

#### *Monitoring on-premises Python applications running on VMs or physical nodes*

You can configure the Python data collector to monitor on-premises Python applications running on VMs or physical nodes. During data collector deployment, the server information must be provided so that the data collector can be configured to connect to the appropriate server. The server information is provided as a configuration package for downloading from the Cloud App Management console.

#### Before you begin

- Ensure that you downloaded the configuration package to obtain the server information. For more information, see [“Obtaining the server configuration information”](#) on page 578.
- Check whether you downloaded the Python data collector package from Passport Advantage. For more information, see [“Downloading the Python data collector”](#) on page 607.

#### About this task

To configure the Python data collector to monitor Python application running on VMs or physical nodes, do the following steps:

#### Procedure

1. Install the Python data collector.

- a) Extract the Data Collector package `ibm_python_datacollector.tgz` with the following command:

```
tar xzf ibm_python_datacollector.tgz
```

- b) Install the data collector package with the following command:

```
pip install --no-index --find-links=/dc_package_extract_folder/python_dc ibm_python_dc
```

Where `dc_package_extract_folder` is the folder where you extract the Python data collector package.

2. Extract the downloaded configuration pack to your python application root path like below command:

```
tar xvf ibm-cloud-apm-dc-configpack.tar
```

You can see a subfolder named `ibm-cloud-apm-dc-configpack` created under the Python application root path. Rename the subfolder to `etc` with the following command:

```
mv ibm-cloud-apm-dc-configpack etc
```

3. From IBM Cloud App Management V2019.4.0, you can use the Python data collector to monitor Python applications running on the web server uWSGI V1.9.0 or later versions. But some configurations are needed to make the Python data collector work well.
  - If you start uWSGI without threads, the threads that are generated by your application and the Python data collector will never run and thus the Python data collector cannot work normally. You must enable uWSGI threads by adding the `--enable-threads` option in the `uwsgi` command. This option is applied automatically when you specify the `--threads` option to configure the number of threads.
  - If you run the `uwsgi` command with the `--master` option, Python scripts and modules are preloaded in the parent master process, and worker processes are forked from the parent master process. In addition, background threads that are created in the master process are killed in worker processes. To make the Python data collector work normally, you must add the `--lazy-apps` option in the `uwsgi` command to use the lazy loading mode.

Command example:

```
uwsgi --enable-threads --master --lazy-apps --processes 4 --http :8002 --wsgi-file=flask_hello.py --callable app
```

**Note:** You can also add the options in a `.ini` configuration file and run the `uwsgi` command with the `.ini` file. For more information, see uWSGI documents.

4. Integrate the installed data collector in your Python application:
  - If your application is based on Django V1.10 or later versions, open `settings.py` of your Django application, and add the following content into first line of section **MIDDLEWARE** in that file:

```
'ibm_python_dc.kpg_dc_django.ResourceMiddleware',
```
  - If your application is based on Django V1.9 or older versions, add the following content into the first line of the section **MIDDLEWARE\_CLASS** in that file:

```
'ibm_python_dc.kpg_dc_django.ResourceMiddleware',
```
  - If your application is based on Flask, add the Python data collector wsgi middleware in your Python application file, for example, if you run `export FLASK_APP=run.py`, then edit the `run.py` file and ensure that the Python data collector wsgi middleware are added in front of other middlewares. Example:

```
from flask import Flask
from flask_restful import Api
from api.board import Article
from api.auth import Login, Register, RefreshToken
from middleware import Test
from werkzeug.middleware.dispatcher import DispatcherMiddleware

api.add_resource(Login, '/login')
api.add_resource(Register, '/register')

from ibm_python_dc.kpg_dc_wsgi import ResourceMiddleware
app.wsgi_app = ResourceMiddleware(app.wsgi_app)
```

```

app.wsgi_app = DispatcherMiddleware(serve_frontend, {
    '/test': test,
    '/admin': admin,
})

.....
.....

```

5. Restart your Python application.

### **Monitoring on-premises Python applications running in individual Docker container**

You can configure the Python data collector to monitor on-premises Python applications running in individual Docker container. During data collector deployment, the server information must be provided so that the data collector can be configured to connect to the appropriate server. The server information is provided as a configuration package for downloading from the Cloud App Management console.

#### **Before you begin**

- Ensure that you downloaded the configuration package to obtain the server information. For more information, see [“Obtaining the server configuration information”](#) on page 578.
- Check whether you downloaded the Python data collector package from Passport Advantage and placed it with Python application Dockerfile in the same directory. For more information, see [“Downloading the Python data collector”](#) on page 607.

#### **About this task**

To configure the Python data collector in individual Docker container mode, do the following steps:

#### **Procedure**

1. Update the Dockerfile to add the following lines to install Python data collector for your Python application. Add the following lines in Dockerfile:

```

ADD ibm-cloud-apm-dc-configpack.tar /application/root/path
RUN mv /application/root/path/ibm-cloud-apm-dc-configpack /application/root/path/etc
ADD ibm_python_datacollector.tar.gz /application/root/path
RUN pip install --no-index --find-links=/application/root/path/python_dc ibm_python_dc

```

2. From IBM Cloud App Management V2019.4.0, you can use the Python data collector to monitor Python applications running on the web server uWSGI V1.9.0 or later versions. But some configurations are needed to make the Python data collector work well.
  - If you start uWSGI without threads, the threads that are generated by your application and the Python data collector will never run and thus the Python data collector cannot work normally. You must enable uWSGI threads by adding the `--enable-threads` option in the **uwsgi** command. This option is applied automatically when you specify the `--threads` option to configure the number of threads.
  - If you run the **uwsgi** command with the `--master` option, Python scripts and modules are preloaded in the parent master process, and worker processes are forked from the parent master process. In addition, background threads that are created in the master process are killed in worker processes. To make the Python data collector work normally, you must add the `--lazy-apps` option in the **uwsgi** command to use the lazy loading mode.

Command example:

```

uwsgi --enable-threads --master --lazy-apps --processes 4 --http :8002 --wsgi-
file=flask_hello.py --callable app

```

**Note:** You can also add the options in a `.ini` configuration file and run the **uwsgi** command with the `.ini` file. For more information, see uWSGI documents.

3. Integrate the installed data collector in your Python application:



- If your application is based on Django V1.10 or later versions, open `settings.py` of your Django application, and add the following content into first line of section **MIDDLEWARE** in that file:

```
'ibm_python_dc.kpg_dc_django.ResourceMiddleware',
```

- If your application is based on Django V1.9 or older versions, add the following content into the first line of the section **MIDDLEWARE\_CLASS** in that file:

```
'ibm_python_dc.kpg_dc_django.ResourceMiddleware',
```

- If your application is based on Flask, add the Python data collector wsgi middleware in your Python application file, for example, if you run **export FLASK\_APP=run.py**, then edit the `run.py` file and ensure that the Python data collector wsgi middleware are added in front of other middlewares. Example:

```
from flask import Flask
from flask_restful import Api
from api.board import Article
from api.auth import Login, Register, RefreshToken
from middleware import Test
from werkzeug.middleware.dispatcher import DispatcherMiddleware

api.add_resource(Login, '/login')
api.add_resource(Register, '/register')

from ibm_python_dc.kpg_dc_wsgi import ResourceMiddleware
app.wsgi_app = ResourceMiddleware(app.wsgi_app)

app.wsgi_app = DispatcherMiddleware(serve_frontend, {
    '/test': test,
    '/admin': admin,
})

.....
.....
```

4. Build the new Docker image.
5. Start your Python application with the new created Docker image.

### Customizing the Python data collector

You can set the variables to change the default behavior of the Python data collector.

#### User-defined environment variables for the Python data collector

For Python monitoring in IBM Cloud Private or OpenShift, set the following variables in the application yaml file.

For on-premises applications, set the following variables as environment variables or in the `config.properties` file.

Table 83. Supported user-defined environment variables for Python monitoring

Variable name	Value	Description
KPG_LOG_LEVEL	<ul style="list-style-type: none"> <li>• DEBUG</li> <li>• ERROR</li> <li>• INFO</li> <li>• Warning</li> </ul>	<ul style="list-style-type: none"> <li>• DEBUG: Only useful debug information is printed in the log, for example, collected data, data that are sent to server, and server response.</li> <li>• ERROR: Only information about exceptions and unexpected situations is printed in the log.</li> <li>• INFO: The summary information about the data collector for the user to know what it is doing is printed in the log.</li> <li>• Warning: Warning information is printed in the log.</li> </ul>
KPG_LOG_TOCONSOLE	True	<p>The log is printed to console and you can see the log by running the command <b>kubectl logs &lt;pod name&gt;:</b>.</p> <p><b>Tip:</b></p> <ul style="list-style-type: none"> <li>• If you do not set KPG_LOG_TOCONSOLE to True, two log files <code>kpg_dc.log</code> and <code>kpg_restclient.log</code> are created to record the log messages in a sub folder <code>kpg_logs</code> that is in your Python application root path.</li> <li>• For Python runtime, standard output(<code>stdout</code>) and standard error(<code>stderr</code>) are buffered. If messages are written into <code>stdout</code> or <code>stderr</code>, they might not be flushed in time if buffer is not full. In this case, when exceptions are reported from Python runtime or Django framework, you might not find any information with the command <code>kubectl logs</code> or the log files <code>kpg_dc.log</code> and <code>kpg_restclient.log</code>. To solve this issue, you can disable <code>stdout</code> and <code>stderr</code> buffering by setting environment <code>PYTHONUNBUFFERED</code> to <code>1</code> in the application <code>yaml</code> file.</li> </ul>

Table 83. Supported user-defined environment variables for Python monitoring (continued)

Variable name	Value	Description
KPG_GC_STATS	True	<p>All statistical functions of python garbage collection are enabled. When you set this value to True, it equals running the following command:</p> <ul style="list-style-type: none"> <li>For Python 2.7,           <pre data-bbox="971 457 1339 577">gc.set_debug(gc.DEBUG_STATS   gc.DEBUG_COLLECTABLE   gc.DEBUG_UNCOLLECTABLE   gc.DEBUG_INSTANCES   gc.DEBUG_OBJECTS )</pre> </li> <li>For Python 3.6.x and 3.7.x,           <pre data-bbox="971 682 1339 756">gc.set_debug(gc.DEBUG_STATS   gc.DEBUG_COLLECTABLE   gc.DEBUG_UNCOLLECTABLE</pre> </li> </ul> <p>To disable KPG_GC_STATS, delete this environment variable. Do not set it to False.</p> <p><b>Note:</b> Never set KPG_SAVE_ALL=True in your formal production environment. It is only for the debug mode. Make sure that enough memory is assigned to the application.</p>
KPG_SAVE_ALL	True	<p>All unreferenced objects are saved into <code>gc.garbage</code>, and you must clear <code>gc.garbage</code> every minute (the data collector clears it for you). When the value is set to True, it equals running the following command:</p> <pre data-bbox="971 1266 1266 1297">gc.set_debug(gc.SAVE_ALL)</pre> <p>To disable KPG_SAVE_ALL, delete this environment variable. Do not set it to False.</p> <p><b>Note:</b> Never set KPG_SAVE_ALL=True in your formal production environment. It is only for the debug mode. Make sure that enough memory is assigned to the application.</p>
KPG_ENABLE_OPENTT	"False"	<p>By default, the Python data collector enables OpenTracing function. You can disable OpenTracing by setting the environment variable to false.</p> <p>Example:</p> <pre data-bbox="971 1795 1266 1843">- name: KPG_ENABLE_OPENTT   value: "false"</pre>

Table 83. Supported user-defined environment variables for Python monitoring (continued)

Variable name	Value	Description
OpenTracing sampling: <ul style="list-style-type: none"> <li>• JAEGER_SAMPLER_TYPE</li> <li>• JAEGER_SAMPLER_PARAM</li> </ul>	The default sampler type is "probabilistic", and the default sampler param is 0.01, which means that 1 in 100 traces will be sampled. You can set it to other values. For more information, see <a href="#">Sampling</a> .	When the OpenTracing function is enabled, you can set the OpenTracing sampler type and param. Example: <pre>- name: JAEGER_SAMPLER_TYPE   value: "probabilistic" - name: JAEGER_SAMPLER_PARAM   value: "0.1"</pre>
LATENCY_SAMPLER_PARAM	Any value between 0 and 1	The default value is 0.1, which means getting 1 request out of 10 requests. The value must be between 0 and 1. The value of 0 means no latency data will be collected. The value of 1 means no sampler and all requests data will be collected. Example: <pre>- name: LATENCY_SAMPLER_PARAM   value: "0.2"</pre>

### Uninstalling the Python data collector

To uninstall the Python data collector, roll back the changes that you have made to your application and then update the application deployment.

#### Procedure

1. Remove relative content from your Python application.

- If your Python application is based on Django 1.10 or later versions, remove the following content from section **MIDDLEWARE** in `setting.py` of your Django application:

```
'ibm_python_dc.kpg_plugin.ResourceMiddleware',
```

- If your Python application is based on Django 1.9 or older versions, remove the following content from section **MIDDLEWARE\_CLASS** in `setting.py` of your Django application:

```
'ibm_python_dc.kpg_dc_django.ResourceMiddleware',
```

- If your Python application is based on Flask, remove the Python data collector wsgi middleware content from the application file. For more information, see [“Monitoring on-premises Python applications running in individual Docker container”](#) on page 612.

2. If you deploy your application in individual Docker container or Kubernetes environments, remove all relative content from Dockerfile to not install python data collector and rebuild the Docker image.
3. If you deploy your application on Kubernetes environments, remove all relative content from Python application yml file.
4. If you deploy your application on bare-metal node or VMs, execute the following commands to remove installed Python data collector packages:

```
pip uninstall ibm-python-restclient
pip uninstall ibm-python-dc
```

5. If you configure the Python data collector by creating secret on Kubernetes environments, execute command like below to remove the secret:

```
kubectl -n my_namespace delete secret icam-server-secret
```

6. Restart your application.

## Configuring Ruby application monitoring

The Ruby data collector can provide you with visibility and control of the Ruby application, and help you ensure optimal performance and efficient use of resources.

The Ruby data collector helps you to manage the performance and availability of the following:

- Ruby on Rails application in Kubernetes environments
- Local Ruby on Rails application

For the detailed system requirements of Ruby data collector, see [system requirements for Ruby data collector](#).

### Downloading the Ruby data collector

You can download the Ruby data collector package from Passport Advantage.

#### About this task

To download the Ruby data collector package, complete the following steps:

#### Procedure

1. Review the part numbers and components to download. For more information, see [Part numbers](#).
2. Extract the package to get the Ruby data collector (`ruby_datacollector.tgz`) package file and then get `stacktracer-19.12.00.gem` by running the following command:

```
tar xzf appMgtDataCollectors_2019.4.0.tar.gz
tar xzf app_mgmt_runtime_dc_2019.4.0.tar.gz
tar xzf app_mgmt_runtime_dc_2019.4.0/ruby_datacollector.tgz
```

### Monitoring Ruby applications in Kubernetes environment

Before you monitor Ruby applications in Kubernetes environment, you must connect the data collector to the server by creating a secret. Then, you update your application deployment to monitor the Ruby applications.

#### Before you begin

- Check whether your service account has access to Kubernetes resources. For more information, see [“Authorizing the data collector to access Kubernetes resources” on page 577](#).
- Ensure that you downloaded the configuration package to obtain the server information. For more information, see [“Obtaining the server configuration information” on page 578](#).
- Check whether you downloaded the Ruby data collector package from Passport Advantage. For more information, see [“Downloading the Ruby data collector” on page 617](#).

#### About this task

You can create a secret for the `global.environment` file and the keyfiles that are extracted from the Cloud App Management configuration package. Then, you mount this secret when you deploy the application as a Kubernetes deployment. To enable the Ruby data collector, you need to install `stacktracer-19.12.00.gem` in Ruby application.

#### Procedure

1. Go to the `ibm-cloud-apm-dc-configpack` directory where you extract the configuration package in [“Obtaining the server configuration information” on page 578](#), and run the following command to create a secret to connect to the server, for example, name it as `icam-server-secret`.

```
kubectl -n my_namespace create secret generic icam-server-secret \
--from-file=keyfiles/keyfile.p12 \
--from-file=keyfiles/ca.pem \
--from-file=global.environment
```

Where *my\_namespace* is the namespace where you want to create the secret. If you want to create the secret in the default namespace, remove `-n my_namespace` from the command.

2. Put `stacktracer-19.12.00.gem` into the Ruby application folder, for example, *app\_folder/RubyDC*.
3. In the Dockerfile, insert the following lines between the section of `bundle package` and `bundle install`.

```
RUN cp <app_folder>/RubyDC/stacktracer-19.12.00.gem <app_folder>/vendor/cache/
RUN echo "gem 'zipkin-tracer', '=0.32.4'" >> Gemfile
RUN echo "gem 'stacktracer', '=19.12.00'" >> Gemfile
```

4. Rebuild the application docker image.

```
docker build --network=host -t <image-name>:<version>
```

5. Repush it to a docker repository.

```
docker tag <image-name>:<version> <docker-repository>/<namespace>/<image-name>:<version>
docker push <docker-repository>/<namespace>/<image-name>:<version>
```

6. Redeploy the application deployment in Kubernetes environment, or relaunch the docker image into container.

### Monitoring on-premises Ruby applications

You can configure the Ruby data collector to monitor the on-premises Ruby applications and then send monitoring data to the Cloud App Management server.

#### Before you begin

- Ensure that you downloaded the configuration package to obtain the server information. For more information, see [“Obtaining the server configuration information”](#) on page 578.
- Check whether you downloaded the Ruby data collector package. For more information, see [“Downloading the Ruby data collector”](#) on page 617.

#### Procedure

1. Run `gem install stacktracer-19.12.00.gem` to install the Ruby data collector.
2. Run `echo "gem 'zipkin-tracer', '=0.32.4'" >> Gemfile`
3. Run `echo "gem 'stacktracer', '=19.12.00'" >> Gemfile`
4. Rerun the Ruby on Rails application by running `rails s -p 3000 -b 0.0.0.0`.

#### Customizing the Ruby data collector

You can set the variables to change the default behavior of the Ruby data collector.

#### User-defined environment variables for the Ruby data collector

For Ruby monitoring in Kubernetes environments, set the following variables in the application yml file.

For on-premises applications, set the following variables as environment variables or in the `config.properties` file.

Table 84. Supported user-defined environment variables for Ruby monitoring

Variable name	Value	Description
OPENTRACING_ENABLED	False	By default, the OpenTracing function is enabled. You can disable OpenTracing by setting the environment variable to false.  Example:  <pre>- name: OPENTRACING_ENABLED   value: false</pre>
OPENTRACING_SAMPLER_PARAM	The default value is 0.01, which means that 1 in 100 traces will be sampled. You can set it to other values.	When the OpenTracing function is enabled, you can set the OpenTracing sampler rate. Example:  <pre>- name: OPENTRACING_SAMPLER_PARAM   value: "0.1"</pre>
LATENCY_SAMPLER_PARAM	Any value between 0 and 1	The default value is 0.1, which means getting 1 request out of 10 requests. The value must be between 0 and 1. The value of 0 means no latency data will be collected. The value of 1 means no sampler and all requests data will be collected.  Example:  <pre>- name: LATENCY_SAMPLER_PARAM   value: "0.2"</pre>

### Uninstalling the Ruby data collector

To uninstall the Ruby data collector, roll back the changes that you have made to your application and then update the application deployment.

#### Procedure

- To uninstall `stacktracer-19.12.00.gem` from Ruby application in Kubernetes environment, do the following steps:

- Remove the following lines in Docker file:

```
RUN cp <app_folder>/RubyDC/stacktracer-19.12.00.gem <app_folder>/vendor/cache/
RUN echo "gem 'stacktracer', '=19.12.00'" >> Gemfile
```

- Remove `stacktracer-19.12.00.gem` from the Ruby application folder, for example, `rm -f <app_folder>/RubyDC/stacktracer-19.12.00.gem`.
- Rebuild the application docker image.

```
docker build --network=host -t <image-name>:<version>
```

- Repush it to a docker repository.

```
docker tag <image-name>:<version> <docker-repository>/<namespace>/<image-name>:<version>
docker push <docker-repository>/<namespace>/<image-name>:<version>
```

- Redeploy the application deployment in Kubernetes environment, or relaunch the docker image into container.
- To uninstall the Ruby data collector in on-premises environment, do the following steps:
    - Run `gem uninstall stacktracer`.

- b) Remove gem 'stacktracer', '=19.12.00' from the Gemfile.
- c) Rerun the Ruby On Rails application by `rails s -p 3000 -b 0.0.0.0`.

## Synthetics PoP

Use the Synthetics PoP to monitor your REST calls and other urls from multiple locations. Create synthetic tests and schedule them to run on a predefined schedule. Monitor both the availability and response time of you websites.

Follow the steps to install the Synthetics PoP, create synthetic tests and view test results.

1. Install Synthetics PoP. After you install the Cloud App Management server, you can install and configure a Synthetics PoP, which enables you to create and run synthetic tests. For more information, see [“Installing Synthetics PoP”](#) on page 620.
2. Navigate to Synthetics page. Click **Administration > Synthetics Configure** on the Cloud App Management console to access the **Synthetics** page with a list of existing tests. Click the **Create** button to open the **Create synthetic test** page, where you can create a new synthetic test.
3. Create synthetic tests based on your requirement. There are four types of synthetic tests:
  - REST API test to test and monitor your REST calls and other urls in response to REST calls. For more information, see [“Creating a REST API test”](#) on page 623.
  - Scripting REST API synthetic test to test and monitor a number of REST APIs in a sequence. For more information, see [“Creating a Scripting REST API synthetic test”](#) on page 627.
  - Web page synthetic test to test a single web page for availability and browser response time. For more information, see [“Creating a web page synthetic test”](#) on page 632.
  - Selenium script test to test simulated user behavior. For more information, see [“Creating a Selenium script test”](#) on page 633.
4. View synthetic test results. Use the Synthetic results tab on the Cloud App Management console to visualize your synthetic tests. For more information, see [“Viewing Synthetic test results”](#) on page 641.

## Installing Synthetics PoP

After you install Cloud App Management server, you can install and configure a Synthetics PoP, which enables you to create and run synthetic tests.

### About this task

There are two components to download:

- Download the data collectors eImage installation tar file `appMgtDataCollectors_2019.4.0.2.tar.gz` from Passport Advantage. This contains a sub package that is called `app_mgmt_syntheticpop_xlinux.tar.gz`, which is the Synthetics PoP installation media.
- Download the data collector configuration package from the Cloud App Management console, use this configuration package (ConfigPack) to configure the Synthetics PoP. The ConfigPack contains the ingress URLs and authentication information that is required to configure the Synthetics PoP to communicate with the Cloud App Management server.

**Note:** For upgrade steps, see [“Upgrading the Synthetics PoP server”](#) on page 1393.

### Procedure

1. Download and unpack the data collectors installation eImage from Passport Advantage (`appMgtDataCollectors_2019.4.0.2.tar.gz`). You will see the Synthetics PoP `app_mgmt_syntheticpop_xlinux.tar.gz` installation file. See [“Downloading agents and data collectors from Passport Advantage”](#) on page 194.
2. Download the data collector configuration package:



- a) Log in to the Cloud App Management console, click the **Get Started** link on the Welcome page, then select **Administration > Integrations**.
  - b) Click the **New integration** button.
  - c) In the **Standard monitoring agents** section, select the **ICAM Data Collectors Configure** button.
  - d) Select **Download file** and specify the directory where you want to save the compressed data collector configuration package, `ibm-cloud-apm-dc-configpack.tar`.
3. Move the extracted `app_mgmt_syntheticpop_xlinux.tar.gz` file and the downloaded ConfigPack to the VM where you want to deploy the Synthetics PoP.

Example that uses secure copy:

```
scp my_path_to_download/appMgtDataCollectors_2019.4.0.2.tar.gz
root@my.env.com:/my_path_to_destination
scp my_path_to_download/ibm-cloud-apm-dc-configpack.tar root@my.env.com:
/my_path_to_destination
```

where

`my_path_to_download` is the path to where the installation tar file or configuration package file was downloaded.

`root` is your user ID on the destination host `my.env.com` where Synthetics PoP is to be installed.

`my_path_to_destination` is the path on the destination host where you want to install Synthetics PoP.

4. Make sure that the prerequisites are met to install the Synthetics PoP in your environment, run the following script.

```
./precheck.sh
```

5. Go to the unpacked folder and run the following command to configure the Synthetics PoP by specifying the downloaded ConfigPack file.

```
./config-pop.sh -f ibm-cloud-apm-dc-configpack.tar
```

You are prompted to enter the following parameters:

Table 85. Synthetics PoP parameters

Property	Required/Optional	Comment
Name	Required	This is displayed as the <b>Locations</b> value in the Synthetics tests tab in the Cloud App Management console. This value cannot be changed if you rerun the configuration.
Country	Required	
City	Required	
Description	Optional	
Agent proxy server	Optional	This is the proxy server address (ip:port) for communicating with the Cloud App Management server.

Table 85. Synthetics PoP parameters (continued)

Property	Required/Optional	Comment
Playback proxy mode	Optional	This is the proxy type for communicating with the web application that is being monitored. This is displayed as the <b>URL</b> in the Cloud App Management console. The value can be:  <b>no</b> Use no proxy. <b>manual</b> Configure the proxy with a proxy server IP address, and port number. <b>pac</b> Use automatic proxy configuration.
Playback Proxy server	Optional	This is the <b>ip:port</b> value for a manual proxy.
Playback Proxy bypass list	Optional	This is the ignore list for a manual proxy.
Playback Proxy pac URL	Optional	This is the URL for a pac proxy.

Validate the deployment:

- When the installation script is complete, run the following command to start the Synthetics PoP:

```
./start-pop.sh
```

- Optional: Other activities that you can complete are:

#### Update the Synthetics PoP

Rerun `./config-pop.sh -f ibm-cloud-apm-dc-configpack.tar` as described in step “5” on page 621. The Synthetics PoP name cannot be changed.

#### Stop the Synthetics PoP

Run `./stop-pop.sh`

#### Delete the Synthetics PoP

Stop the Synthetics PoP and in the **Locations** section in the Synthetic test editor page, click the **delete** button.

#### Collect logs

Run the following command to collect the log files: `./pdcollect.sh`

- Optional: From IBM Cloud App Management V2019.4.0, you can choose to store the playback data in the Synthetics PoP side in case of any connection issues between the Synthetics PoP and the Cloud App Management server. The playback data is first stored in memory and if its size exceeds the maximum caching size, which is set by the **CACHE\_REDIS\_MAX\_SIZE\_MB** property, the playback data can be persisted on the disk in the docker container. The Synthetics PoP can wait and retry posting all the stored playback data to the Cloud App Management server until it can connect to the server.

By default this function is disabled and thus the playback data is lost when the Synthetics PoP cannot connect to the Cloud App Management server.

- To enable this function, you need to designate the maximum caching size for the playback data. Go to the Synthetics PoP installation path and set the **CACHE\_REDIS\_MAX\_SIZE\_MB** property in the `openpop.properties` file. For example:

```
CACHE_REDIS_MAX_SIZE_MB=2048
```

In this example, 2 GB playback data can be stored in memory at most and the other playback data can be stored on the disk. You can change the value as you need based on the memory capacity in your environment.

**Note:** Because some other processes are also using the memory, it is not suggested to set the **CACHE\_REDIS\_MAX\_SIZE\_MB** property too large.

- To disable this function, set the **CACHE\_REDIS\_MAX\_SIZE\_MB** property in the `openpop.properties` file to 0.

```
CACHE_REDIS_MAX_SIZE_MB=0
```

**Note:** After the **CACHE\_REDIS\_MAX\_SIZE\_MB** is set to 0, the playback data cannot be stored in both memory and the disk when a connection issue happens between the Synthetics PoP and the Cloud App Management server.

## Results

You can now log in to Cloud App Management console and create a synthetic test. The Synthetics PoP is shown as a Location based on the name value you specified. For how to create different types of synthetic tests, see [“Synthetics PoP” on page 620](#).

## Creating a REST API test

Create a REST API test to monitor the availability and performance of your web application and other URLs in response to REST calls.

### About this task

Create a REST API test to test the response time and availability of your web applications by using the following HTTP methods: GET, POST, PUT, and DELETE.

### Procedure

Name and description

1. Enter a meaningful name for your test in the **Name** field. Add a description of the purpose of your test to the **Description** field.
2. In the **Test type** section, select **REST API**.

Request

3. In the **Request** section, select the type of method from the **Method** list and enter a URL that you want to test with this method. You can choose GET, PUT, POST, or DELETE. If you choose the PUT or POST method, you can enter body content to test in the **Request body (optional)** field.
4. Optional: You can configure your test to include a particular header and value. Enter a header name and header value in the **Header** fields. If the web application that you want to test requires a user login and password, enter **Authorization** into the **Header** field. Enter the word **Basic**, a space character, and the base64 encoded value of your **username:password** into the **Header value** field.

For example, if your username is Aladdin and your password is OpenSesame, then enter the word **Basic**, a space character, and the base64 encoded value for Aladdin:OpenSesame into the Header value field.

Response validation

5. Configure the warning and critical events conditions for your synthetic test in the **Response Validation** section. You can see two conditions based on response time are provided to trigger events. By default, a response time over 5 seconds triggers a warning event and a response time over

10 seconds triggers a critical event. You can change the response time in the **Threshold Value** field or change the unit to milliseconds or seconds in the **Unit** field for each condition. Response times that exceed threshold values in your warning and critical conditions trigger events.

Further customization of warning and critical events can be done in the next configuration stage. For more information, see Event triggers later in this procedure.

For more detail in relation to event triggers default behavior and how event triggers function across multiple Synthetics PoP locations, see [“Event generation” on page 639](#).

6. Optional: You can click **Add Condition** to define and add customized response validation conditions. Customized response validation conditions are evaluated in aggregate to generate an event. Each test can generate up to a total of three alerts. Your test reports the event with the highest severity until all conditions that cause events are resolved.

- a) Select one of the following options from the **Validate** drop-down menu:

**Header response code**

Select **Header response code** to test for one or for a range of HTTP response codes.

**Header property**

Select **Header property** to test for a particular HTTP header field property and value.

**Body JSON**

Select **Body JSON** to test for a particular property from a JSON body.

- b) Enter a property to test for in the **Target** field and a value to test for in the **Threshold Value** field.

**Note:**

- For **Header response code** conditions, the **Target** field is fixed to Response code. You need to enter a numerical value to the **Threshold Value** field and the value can't be greater than 600.
- For **Header property** conditions, you need to enter no more than 100 characters for both the **Target** and the **Threshold Value** fields.
- For **Body JSON** conditions, you need to enter no more than 150 characters for the **Target** field and no more than 100 characters for the **Threshold Value** field.

- c) Select an operation from the **Operation** drop-down menu.

**Note:**

- The available operations for **Header response code** conditions are: =, >, <, <=, =, !=
- The available operations for **Header property** conditions are: =, !=, contains, doesn't contain.
- The available operations for **Body JSON** conditions are: =, >, <, <=, =, !=, contains, doesn't contain.

- d) Choose an **Event severity** of **Warning** or **Critical** for your condition. For examples, see [“Customizing response validation” on page 625](#).

Verify

7. Click **Verify** to determine whether your test request is valid. No response validation takes place during test verification. Your validated test is displayed in the verified test window. You can rename or delete your test in the verified test window. Click **Next**.

Review & Finish

8. Enter an Interval and Testing frequency.

**Interval**

Defines how often the test runs in minutes or hours.

**Testing frequency**

Determines whether your test runs from all locations simultaneously or from a different location at each interval. Select **Simultaneous** to run your test from all locations simultaneously, or select **Staggered** to run your test from a different selected location at each interval.

9. The **Locations** sections lists the Synthetics PoP that are installed. The first Synthetics PoP is selected by default. You can run your test from one or more synthetic pop servers.

Select the synthetic pop servers where you want your synthetic test to run. To create a new Location, see [“Installing Synthetics PoP”](#) on page 620.

#### Event triggers

10. By default a critical alert is triggered if a synthetic test playback fails (returns a code 400 or above). To stop this behavior, set **Trigger an event if a failure is detected** to **Off**. To increase the number of failures allowed before a critical alert is triggered, change the value between **Trigger an event if the test fails** and **consecutive times** under the **Failure** section. The default number of consecutive failures is 0.

By default, a critical event is triggered if a synthetic test playback response time is >10 seconds. By default, a warning event is triggered if a synthetic test playback response time is >5 seconds.

To increase the number of slow response times that must occur before a critical or warning event is triggered, change the value between **Trigger an event if a threshold is breached** and **consecutive times** under the **Slow response threshold** section. The default number of slow response times is 0.

For more detail in relation to event triggers default behavior and how event triggers function across multiple Synthetics PoP locations, see [“Event generation”](#) on page 639.

#### Customizing response validation

Add customized response validation conditions based on; header response code, header property, or body JSON values.

#### Conditions based on header response code

When you create a condition that uses the **header response code** in the **Validate** field, set the **Target** field by using the following example:

header response code	response code	>=	302	warning
----------------------	---------------	----	-----	---------

In this example, a 302 code indicates a redirect. If a redirect occurs, in the **Synthetic results** tab, in the **Test instance breakdown** widget, you see an error, click the error to see a breakout similar to this:

Timestamp	2019-04-10T18:31:38.289Z
Alarm condition	statuscode equals 302
Event type	Verification
Additional details	Warning: alarm condition satisfied!
Operator	equals
Target	statuscode

#### Conditions based on header property

When you create a condition that uses the **header property** in the **Validate** field, set the **Target** field by using one of the following examples:

Example 1:

header property	content-encoding	=	"gzip"	warning
-----------------	------------------	---	--------	---------

Response header sample:

```

accept-ranges: bytes
cache-control: max-age=301
content-encoding: gzip
content-length: 9168
content-security-policy: upgrade-insecure-requests

```

Example 2:

header property	content-type	contains	application/json	warning
-----------------	--------------	----------	------------------	---------

Response header sample:

```

date: Wed, 13 Mar 2019 01:28:46 GMT
content-type: application/json; charset=utf-8
x-powered-by: Express
access-control-allow-origin: *
etag: W/"1bb-D+c3sZ5g5u/nmLPQR11uVo2heAo"
cf-ray: 4b6a3bac0a41cc2a-SIN
content-encoding: br
X-Firefox-Spdy: h2

```

### Conditions based on body JSON

When you create a condition that uses **body json** in the **Validate** field, set the **Target** field by using one of the following examples, which is based on the JSON code sample:

Set **Target** to a string or number property in a JSON object, for example:

body json	page	>	2	warning
-----------	------	---	---	---------

Set **Target** to a string or number property in an object within a JSON. Syntax is: *parent\_object\_name.property\_name*, for example:

body json	properties.key.1	>	10	warning
-----------	------------------	---	----	---------

Set **Target** to a string or number property in an array within a JSON object. An element of an array can be accessed by using *[item\_index]*, for example:

body json	date[0].first_name	contains	"George"	warning
-----------	--------------------	----------	----------	---------

#### Note:

1. Wildcards are not supported for property, you must use an absolute string or number only.
2. Numerical values that you enter in the **Value** field are treated as numbers and not strings by default. Use quotation marks "" to distinguish between a string and a number. For example, to test for the string 123, enter "123" in the **Value** field. To check for the number 400, enter 400 without any quotation marks.

JSON Code sample:

```

{
  "page": 1,
  "per_page": 3,
  "total": 12,
  "total_pages": 4,
  "properties": {
    "key1": 10,
    "key2": "testkey3"
  },
  "data": [
    {
      "id": 1,
      "first_name": "George",
      "last_name": "Bluth",
    }
  ]
}

```

```
    "id": 2,  
    "first_name": "Janet",  
    "last_name": "Weaver",  
  },  
  {  
    "id": 3,  
    "first_name": "Emma",  
    "last_name": "Wong",  
  }  
}]  
}
```

## Creating a Scripting REST API synthetic test

Use a scripting REST API test to test a sequence of REST APIs. Use a node . js script to test your sequenced REST APIs.

### Procedure

Name and Description

1. Enter a meaningful name for your test in the **Name** field. Add a description of the purpose of your test to the **Description** field.

Test type

2. Select Scripting REST API.

Request

3. You can choose one of the following options from the **Upload options** list to provide the test script:

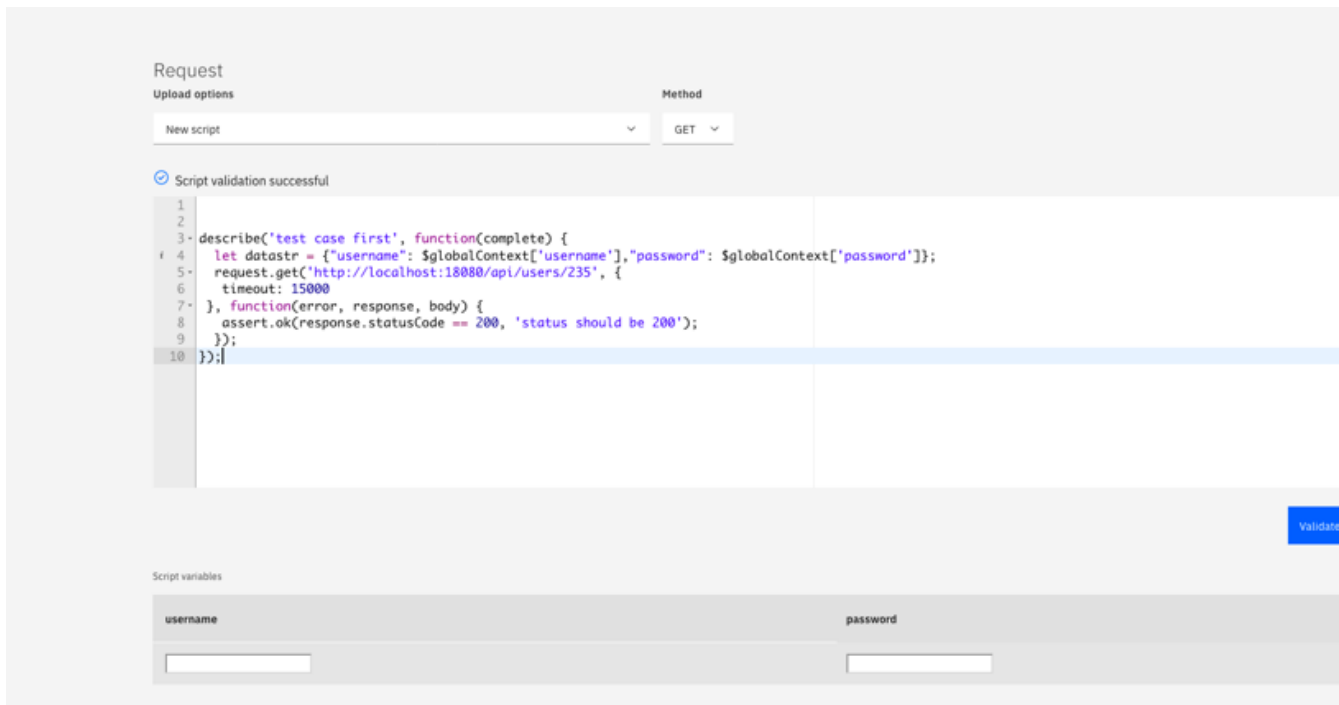
- **Upload a JS file**

Use the **Upload a JS file** option to upload a node . js test script file. For information about how to create a node . js script to test, see [“Create a REST API test case” on page 629](#).

- **New script**

Use the **New script** option to create a new test script. Type your script in the script editor. Choose a method from the **Method** list. Then the template script for the method that you choose is added to the script editor. You only need to edit the template script to meet your goal, instead of typing the script manually.

If you use **\$globalContext[VAR\_NAME]** to pass variables in the script, you can see the **Script variables** section after clicking **Validate**. You can type the variable values in the **Script variables** section. The **Script variables** section is hidden by default.



**Note:** You can verify whether your script is grammatically right by clicking **Validate**. You can also download the script to your local computer by clicking **Download**.

- **Add a template**

This option works in the same way as the **New script** option and is designed to be removed from Cloud App Management console **V2020.1.0**.

#### Response validation

4. Configure the warning and critical events conditions for your synthetic test in the **Response Validation** section. You can see two conditions based on response time are provided to trigger events. By default, a response time over 5 seconds triggers a warning event and a response time over 10 seconds triggers a critical event. You can change the response time in the **Threshold Value** field or change the unit to milliseconds or seconds in the **Unit** field for each condition. Response times that exceed threshold values in your warning and critical conditions trigger events.

Further customization of warning and critical events can be done in the next configuration stage. For more information, see [Event triggers](#) later in this procedure.

For more detail in relation to event triggers default behavior and how event triggers function across multiple Synthetics PoP locations, see [“Event generation” on page 639](#).

#### Verify

5. Click **Verify** to determine whether your test request is valid. No response validation takes place during test verification. Your validated test is displayed in the verified test window. You can rename or delete your test in the verified test window. Click **Next**.

#### Review and Finish

6. Enter an Interval and Testing frequency.

##### **Interval**

Defines how often the test runs in minutes or hours.

##### **Testing frequency**

Determines whether your test runs from all locations simultaneously or from a different location at each interval. Select **Simultaneous** to run your test from all locations simultaneously, or select **Staggered** to run your test from a different selected location at each interval.

#### Locations

7. The **Locations** sections lists the Synthetics PoP that are installed. The first Synthetics PoP is selected by default. You can run your test from one or more synthetic pop servers.



Select the synthetic pop servers where you want your synthetic test to run. To create a new Location, see [“Installing Synthetics PoP” on page 620](#).

#### Script variables

8. If you introduced any variables in the `node.js` script, they are requested at this point.

#### Event triggers

9. By default a critical alert is triggered if a synthetic test playback fails (returns a code 400 or above). To stop this behavior, set **Trigger an event if a failure is detected** to **Off**. To increase the number of failures allowed before a critical alert is triggered, change the value between **Trigger an event if the test fails** and **consecutive times** under the **Failure** section. The default number of consecutive failures is 0.

By default, a critical event is triggered if a synthetic test playback response time is >10 seconds. By default, a warning event is triggered if a synthetic test playback response time is >5 seconds.

To increase the number of slow response times that must occur before a critical or warning event is triggered, change the value between **Trigger an event if a threshold is breached** and **consecutive times** under the **Slow response threshold** section. The default number of slow response times is 0.

For more detail in relation to event triggers default behavior and how event triggers function across multiple Synthetics PoP locations, see [“Event generation” on page 639](#).

### Create a REST API test case

Create a REST API test case with Java script. Upload this script to Cloud App Management and run it on a schedule.

Create and schedule a Java script to test your REST APIs. Use Java script to:

- Monitor your REST API in every stage of DevOps pipeline.
- Automate detection of REST API defects in the continuous pipeline.
- Test REST API transaction response time in the continuous pipeline.
- Test REST API uptime and transaction business logic in production.

### How to write a REST API test case with Java script

1. Use `describe` to create a script step.

Every request should be in a step.

Every step should have only one request, use multiple steps if there are multiple requests.

Steps run in sequence. Subsequent steps run after the previous step is finish and the complete call is reached.

The syntax is as follows:

```
describe(stepName, function(complete){}) complete
```

2. To pass variables from the synthetic test configuration, use the following syntax `$globalContext[VAR_NAME]`

For example,

```
describe('step 1', function(complete){
  let datastr = {"name": $globalContext['name'], "job": $globalContext['job']};
  request.post("http://localhost:18080/api/users",
    {headers: {'Content-type': 'application/json'}, body: JSON.stringify(datastr)
  },
  function(error, response, body){
    assert.ok (response && response.statusCode == 200, "Response code is not 200");
    var bodyObj = JSON.parse(body);
    assert.ok(bodyObj['name'] == $globalContext['name'], "Pass parameter name failed");
    assert.ok(bodyObj['job'] == $globalContext['job'], "Pass parameter jobfailed");
  }
})
```

```
);  
});
```

3. To pass data to the next step or next N step use the following syntax: `complete(DATA)`. Where DATA can be any type, for example, string, number or object. For example:

```
complete(['value', 1234, 'value2']), complete({"key1":"value1","key2":"value2"})
```

Pass data to next steps from step 1, for example:

```
complete({"key1":"value1","key2":"value2"})
```

Read data from previous step by using the following syntax: `$stepContext['preStepResult']`  
For example, Read data "key1" in step 2 by

```
var k = $stepContext['preStepResult']['key1']
```

Read "key1" in step 3 by, for example:

```
var k = $stepContext['key1']
```

If you do not need to pass anything to the next step, use `complete()`.

4. To validate your results, call an `assert` method to validate the endpoint response. If the condition of the assert method are met, an alert notification is triggered.

For example:

```
assert.ok(response && response.statusCode == 200, "step failed ,error is " + error);
```

Validate response content, for example:

```
let bodyObj = typeof body == 'string': JSON.parse(body): body;  
assert.ok(bodyObj.id != null, "Resource create failed, not resource ID");
```

5. Use the following syntax for a GET request:

```
let baseUrl = 'https://reqres.in/api/messages';  
// Make GET request  
describe('get messages', function (complete) {  
  request.get(baseUrl, {}, function (error, response) {  
    // Validate the response code, if assertion fails, log "failed to get message "  
    plus results as error message on synthetic dashboard  
    assert.ok(response && response.statusCode == 200, "failed to get message" + error);  
    complete();  
  });  
});
```

6. Use the following syntax for a POST form request:

```
let baseUrl = 'http://localhost:3000';  
describe('post form content',function(complete){  
  // Define endpoint URL  
  let url = baseUrl + "/messages";  
  // Define form content data  
  let formData = {"text": "Hi Meg is here"};  
  // Define headers such as "context-type"  
  let header1 = {"context-type": "application/json"};  
  // Make POST request  
  request.post(url, {header: header1, form: formData}, function(error, response) {  
    // Validate the response code, if assertion fails, log "failed to create message, error is "  
    plus results as error message on synthetic dashboard  
    assert.ok(response && response.statusCode == 200, "failed to create message, error is " +  
    error);  
    complete();  
  });  
});
```

7. Use the following syntax for a POST JSON request:

```

//Send Json content
let baseUrl = 'https://reqres.in';
describe('post json content', function (complete) {
  // Define endpoint URL
  let url = baseUrl + "/api/users";
  // Define JSON data
  let data = { "job": "leader", "name": "morpheus" };
  // Define headers
  let header1 = { "contex-type": "application/json" };
  // Make POST request
  request.post(url, { header: header1, json: data }, function (error, response) {
    // Validate the response code, if assertion fails, log "failed to create message,
    error is " plus results as error message on synthetic dashboard
    assert.ok(response && response.statusCode == 201, "failed to create message, error
    is " + error);
    complete();
  });
});

```

8. Use the following syntax to send a TLS/SSL Protocol request with cert:

```

describe('test TLS/SSL',function(complete){
  const cert= <client.crt string>
  , key= <client.key string>
  , ca= <ca.cert.pem string>;

  const options = {
  url: 'https://api.some-server.com/',
  cert: cert,
  key: key,
  passphrase: <password>,
  ca: ca
  };
  request.get(options, function(error, response, body){
  complete()
  });
});

```

9. Use the following syntax to send a basic authentication request:

```

describe('step 1', function(complete){
  request.get('http://some.server.com/', {
  // Define authentication credentials
  'auth': {
  'user': 'username',
  'pass': 'password',
  'sendImmediately': false
  }
  });
});

```

10. Use the following syntax to send a bear authentication request. Set the bearer value in the **auth** parameter. The value can be either a string or a function returning a string.

```

describe('step 1', function(complete){
  request.get('http://some.server.com/', {
  'auth': {
  'bearer': authToken
  }
  });
});

```

11. Use the following syntax to send a request with http proxy:

```

describe('step 1', function(complete){
  request.get('http://www.ibm.com',{
  "proxy": "http://PROXY_HOST:PROXY_PORT"
  }, function (error, response, body){
  //console.log(body);
  });
});

```

12. Optional: Server Name Indication(SNI) is supported from IBM Cloud App Management V2019.3.0. You can add servername in the request to get the right SSL certificate from the specific server. Use the following syntax to send request to SNI enabled web server:

```
describe('step 1',function(complete){
  request.get({
    url: 'https://some.server.name',
    servername: 'some.server.name'
  }, function(error, response, body){
    console.log(response.statusCode);
    complete();
  });
});
```

## Creating a web page synthetic test

Use a Webpage synthetic test to test a single web page.

### About this task

#### Procedure

##### Name and Description

1. Enter a meaningful name for your test in the **Name** field. Add a description of the purpose of your test to the **Description** field.

##### Test type

2. Select Web page.

##### Request

3. Enter the URL of web page.

##### Response validation

4. Configure the warning and critical events conditions for your synthetic test in the **Response Validation** section. You can see two conditions based on response time are provided to trigger events. By default, a response time over 5 seconds triggers a warning event and a response time over 10 seconds triggers a critical event. You can change the response time in the **Threshold Value** field or change the unit to milliseconds or seconds in the **Unit** field for each condition. Response times that exceed threshold values in your warning and critical conditions trigger events.

Further customization of warning and critical events can be done in the next configuration stage. For more information, see Event triggers later in this procedure.

For more detail in relation to event triggers default behavior and how event triggers function across multiple Synthetics PoP locations, see [“Event generation” on page 639](#).

##### Blocking and filtering

5. **Whitelist** and **Blacklist** fields determine which resources you send requests to and contribute to the metrics and status of your service tests. In the **Blacklist** field, enter any URL or domains that you want to block from any requests and metric calculations. In the **Whitelist** field, enter URL or domains that you want to include in metric calculations. Any non matching domains and URL will be blacklisted.

**Note:** Each URL or domain must be 200 characters or fewer. Use commas (,) to separate them and the wildcard symbol (\*) to filter them. For example: `ibm.com,*developerworks*,*.s81c.com/*`  
Up to 20 comma separated entries are allowed in the **Blacklist** field and up to 10 comma separated entries are allowed in the **Whitelist** field.

##### Authentication

6. If the web page you are testing requires authentication (NTLM or basic), enter the username and password.

##### Verify

7. Click **Verify** to determine whether your test request is valid. No response validation takes place during test verification. Your validated test is displayed in the verified test window. You can rename or delete your test in the verified test window. Click **Next**.

#### Review and Finish

8. Enter an Interval and Testing frequency.

##### **Interval**

Defines how often the test runs in minutes or hours.

##### **Testing frequency**

Determines whether your test runs from all locations simultaneously or from a different location at each interval. Select **Simultaneous** to run your test from all locations simultaneously, or select **Staggered** to run your test from a different selected location at each interval.

#### Locations

9. The **Locations** sections lists the Synthetics PoP that are installed. The first Synthetics PoP is selected by default. You can run your test from one or more synthetic pop servers.

Select the synthetic pop servers where you want your synthetic test to run. To create a new Location, see [“Installing Synthetics PoP” on page 620](#).

#### Event triggers

10. By default a critical alert is triggered if a synthetic test playback fails (returns a code 400 or above). To stop this behavior, set **Trigger an event if a failure is detected** to **Off**. To increase the number of failures allowed before a critical alert is triggered, change the value between **Trigger an event if the test fails** and **consecutive times** under the **Failure** section. The default number of consecutive failures is 0.

By default, a critical event is triggered if a synthetic test playback response time is >10 seconds. By default, a warning event is triggered if a synthetic test playback response time is >5 seconds.

To increase the number of slow response times that must occur before a critical or warning event is triggered, change the value between **Trigger an event if a threshold is breached** and **consecutive times** under the **Slow response threshold** section. The default number of slow response times is 0.

For more detail in relation to event triggers default behavior and how event triggers function across multiple Synthetics PoP locations, see [“Event generation” on page 639](#).

## Creating a Selenium script test

Use a Selenium script test if you want to simulate user interactions with your web application. Record a synthetic script by using the Firefox web browser and the Selenium IDE add-on. Record user actions on a web page, such as loading a page, clicking a link, or selecting an object. When Selenium IDE is recording, it generates a command for each user action in a script. Use a Selenium script test to replay the Selenium script at set intervals and at different locations.

### **Procedure**

#### Name and Description

1. Enter a meaningful name for your test in the **Name** field. Add a description of the purpose of your test to the **Description** field.

#### Test type

2. Select Selenium script.

#### Request

3. Upload a Selenium .side file. To record a Selenium .side file, see [“Recording a Selenium script” on page 635](#).
4. If the uploaded Selenium .side file includes variables, you can see the **Script variables** section in Cloud App Management console, where you can input the variable values as you need.

**Note:** By default, the **Script variables** section is hidden. You can see the section only when the uploaded Selenium script uses variables. For more information about Selenium script variables, see [“Passing variable values to Selenium script”](#) on page 638.

#### Response validation

5. Configure the warning and critical events conditions for your synthetic test in the **Response Validation** section. You can see two conditions based on response time are provided to trigger events. By default, a response time over 5 seconds triggers a warning event and a response time over 10 seconds triggers a critical event. You can change the response time in the **Threshold Value** field or change the unit to milliseconds or seconds in the **Unit** field for each condition. Response times that exceed threshold values in your warning and critical conditions trigger events.

Further customization of warning and critical events can be done in the next configuration stage. For more information, see Event triggers later in this procedure.

For more detail in relation to event triggers default behavior and how event triggers function across multiple Synthetics PoP locations, see [“Event generation”](#) on page 639.

#### Blocking and filtering

6. **Whitelist** and **Blacklist** fields determine which resources you send requests to and contribute to the metrics and status of your service tests. In the **Blacklist** field, enter any URL or domains that you want to block from any requests and metric calculations. In the **Whitelist** field, enter URL or domains that you want to include in metric calculations. Any non matching domains and URL will be blacklisted.

**Note:** Each URL or domain must be 200 characters or fewer. Use commas (,) to separate them and the wildcard symbol (\*) to filter them. For example: `ibm.com,*developerworks*,*.s81c.com/*` Up to 20 comma separated entries are allowed in the **Blacklist** field and up to 10 comma separated entries are allowed in the **Whitelist** field.

#### Authentication

7. If the web page you are testing requires authentication (NTLM or basic), enter the username and password.

**Note:** The **Verify test** button will verify the Selenium script only if the Cloud App Management server is running on xLinux platform. If Cloud App Management server is deployed on pLinux or zLinux platform, the **Verify test** button will bypass Selenium script verification.

#### Review and Finish

8. Enter an Interval and Testing frequency.

##### **Interval**

Defines how often the test runs in minutes or hours.

##### **Testing frequency**

Determines whether your test runs from all locations simultaneously or from a different location at each interval. Select **Simultaneous** to run your test from all locations simultaneously, or select **Staggered** to run your test from a different selected location at each interval.

#### Locations

9. The **Locations** sections lists the Synthetics PoP that are installed. The first Synthetics PoP is selected by default. You can run your test from one or more synthetic pop servers.

Select the synthetic pop servers where you want your synthetic test to run. To create a new Location, see [“Installing Synthetics PoP”](#) on page 620.

10. You can customize the Selenium script variables for different locations. Example:

Locations	username	password
mither1_opentt_longrun	demo1	.....
whine1_pop	demo2	.....
wqchrometest	demo3	.....

Where:

- The column **Locations** are the installed Synthetics PoP.
  - The columns **username** and **password** are the variables that are used in the uploaded Selenium script. For more information about Selenium script variables, see [“Passing variable values to Selenium script”](#) on page 638.
11. By default a critical alert is triggered if a synthetic test playback fails (returns a code 400 or above). To stop this behavior, set **Trigger an event if a failure is detected** to **Off**. To increase the number of failures allowed before a critical alert is triggered, change the value between **Trigger an event if the test fails** and **consecutive times** under the **Failure** section. The default number of consecutive failures is 0.

By default, a critical event is triggered if a synthetic test playback response time is >10 seconds. By default, a warning event is triggered if a synthetic test playback response time is >5 seconds.

To increase the number of slow response times that must occur before a critical or warning event is triggered, change the value between **Trigger an event if a threshold is breached** and **consecutive times** under the **Slow response threshold** section. The default number of slow response times is 0.

For more detail in relation to event triggers default behavior and how event triggers function across multiple Synthetics PoP locations, see [“Event generation”](#) on page 639.

### Recording a Selenium script

Use a Selenium test to run a Selenium script that simulates user interactions with your web application. You can create selenium tests to simulate user behavior at your website, at set intervals and at different locations. Record a synthetic script by using the Firefox web browser and the Selenium IDE add-on. With Selenium IDE, you can record user actions on a web page, such as loading a page, clicking a link, or selecting an object. When Selenium IDE is recording, it generates a command for each user action in a script.

### Before you begin

#### You must use the Firefox web browser when recording scripts

Selenium IDE is available only as a Firefox add-on. If Selenium IDE is not installed or running, complete the following steps:

1. Ensure that you are running a version of Firefox 68.0.1 esr or later that supports Selenium IDE 3.12. If you have a later version of Selenium IDE, it is not supported; you must uninstall it and install version 3.12. Turn off automatic updates for Selenium IDE to prevent version upgrades.
2. Download and install Selenium IDE 3.12 from the **Selenium** home page (<https://addons.mozilla.org/firefox/addon/selenium-ide/versions/>). Allow Selenium IDE to install all plug-ins. Restart Firefox.
3. Navigate to the web page that you want to test and close any other tabs. To open Selenium IDE, click **Tools > Selenium IDE**. In the **Selenium IDE** window, ensure that the **Base URL** field contains the URL of the displayed web page. Selenium IDE starts recording all user actions on the displayed web page.

## Selenium .side script format

Scripts created with newer versions of Selenium use the `.side` format. With Selenium IDE 3.12, you can import older scripts that were created with the `.html` format and save to the `.side` format. For more information, see [“Updating scripts from earlier Selenium IDE versions”](#) on page 639.

### About this task

In this task, you perform user actions on a web page and use Selenium IDE to record these actions as commands in a simple script. You can use scripts to monitor the performance and availability of your web application in the Application Performance Dashboard.

### Procedure

Complete the following steps to record a script of user actions on a web page:

1. Click **Record** to start recording a script. Perform user actions on your web page, such as clicking a link.

For every user action on a web page, Selenium IDE records a command and adds it to a script.

For example, complete the following actions to record when a user loads the [IBM Marketplace](#) web page and navigates to a free trial of Cloud APM, in a script:

User action	Commands added to script
To record when the Cloud APM web page on the IBM Marketplace website opens, open the <a href="#">IBM Marketplace</a> web page. Right-click anywhere on the displayed web page and select <b>open</b> .	open
To ensure that the script checks that the web page loads, right-click the title text of the web page (IBM Cloud Application Performance Management) and click <b>Show All Available Commands &gt; verifyTitle IBM Cloud Application Performance Management</b> .	verifyTitle
To record when the user clicks a link to view details about Cloud APM, click the <b>Details</b> link. The <b>Details</b> page loads.	clickAndWait
To ensure that the script checks that the <b>Details</b> page has loaded, right-click on the "Feature spotlights" heading and select <b>Show All Available Commands &gt; verifyText css=h2.heading--TERTIARY</b> .	verifyText
To record when the user clicks a link to view details about how to purchase Cloud APM, click the <b>Purchase</b> link. The <b>Purchase</b> page loads.	clickAndWait
To record when the user clicks a button to register for a free trial of Cloud APM, click the <b>Try Free</b> button.	click

2. In the Selenium IDE window, click **Record** to stop the recording. Click the **Save Project** tool, give your script a meaningful name, and save as a `.side` file (such as `open_webpage.side`).
3. In the Selenium IDE window, review your recorded script. Click the **Table** tab to display the script in a table format. In the Selenium IDE window, click **Play Current Test Case** to test the playback of the script that you recorded.

In this example, Selenium IDE displays the script of user actions on the IBM Marketplace website, as described in step 1.

Command	Target	Value
open	/	



Table 87. Example of a Selenium IDE script recording of user actions on the IBM Marketplace website (continued)

Command	Target	Value
verifyTitle	IBM Cloud Application Performance Management	
clickAndWait	css=ul > #details > a	
verifyText	css=h2.heading--TERTIARY	Feature spotlights
clickAndWait	css=ul > #purchase > a	
click	link=Try Free	

## Results

You recorded a script that you can use to monitor the performance and availability of a web application.

## What to do next

If you recorded a complex script, you can organize your script into simpler scripts, where each script represents a specific business process or user action on your web application.

### Structuring complex scripts

It is good practice to organize complex scripts into separate scripts, where each script represents a typical user or business process that you want to monitor. For example, create separate scripts that record when a user logs in to a website, or searches for an item. If you organize your scripts according to user or business processes, you can then monitor the response of your web application to these specific processes. Organize a complex script into multiple scripts; then, save scripts together in a collection of scripts called a *test suite*.

### About this task

If you create a complex script, you can organize that script into simple scripts that represent different business or user processes on your web application. Save the scripts together as a test suite. You can then use these scripts to monitor the performance and availability of your web application in response to specific user actions. There should be only one test suite and all tests should be added into it.

### Procedure

To organize your complex script into separate scripts, and save your scripts as a test suite, complete the following steps:

1. To create a separate script for each user process that is recorded in your script, click **Tests > +** in Selenium IDE. Give each script a meaningful name that describes the user process and save each script as a `.side` file, such as `load_homepage.side`.  
For more information, see [“Recording a Selenium script” on page 635](#).
2. In Selenium IDE, open a complex script that you recorded previously. Organize your script commands into separate scripts, according to different user actions. **Cut** commands from the original complex script in the **Test Case** window and **Paste** commands into the different **Test Case** window.

For example, the complex script example in [“Recording a Selenium script” on page 635](#) contains Selenium IDE commands for three different user processes.

- Open the Cloud APM home page on the IBM Marketplace website.
- Open the **Details** page on IBM Marketplace.
- Open the **Pricing** page and record when the user opens the registration page for a free trial.

The user actions are then organized into three different scripts.

Table 88. Sample script for opening the IBM Marketplace page ( <i>load_homepage.side</i> )		
Command	Target	Value
open	/	
verifyTitle	IBM Cloud Application Performance Management	

Table 89. Sample script for opening the <b>Details</b> page on IBM Marketplace ( <i>load_products.side</i> )		
Command	Target	Value
clickAndWait	css=ul > #details > a	
verifyText	css=h2.heading--TERTIARY	Feature spotlights

Table 90. Sample script for opening the <b>Purchase</b> and trial registration pages on IBM Marketplace ( <i>load_APM.side</i> )		
Command	Target	Value
clickAndWait	css=ul > #purchase > a	
click	link=Try Free	

3. To put individual test cases into a test suite, change to the **Test suite** window and add tests to the test suite according to the business logic sequence. Finally, click the **Save Project** tool to save the test suite and all tests in the test suite to a *.side* file.

As an example, consider the logical sequence `Load_URL`, `Select Manage inventory`, `Select IBM Machine Type`. When we add these test cases to the test suite, we first check `Load_URL`, followed by `Select Manage inventory`, then `Select IBM Machine Type`

## Results

You recorded a set of scripts that you can use to monitor the performance and availability of your web applications.

## Passing variable values to Selenium script

In some cases, you prefer not to input real values in the Selenium script for security or customizing considerations. For example, you do not want to input the password directly in the script and you want to specify the password later in Cloud App Management console. To achieve this goal, you can store and use variables in the script and specify real values for the variables in Cloud App Management console when you create a Selenium script test. These variable values can be passed from Cloud App Management console to the Selenium script.

## Procedure

1. Record a Selenium script test. See [“Creating a Selenium script test”](#) on page 633.
2. If you want to use variables for some commands instead of using real values, you need to change the script in Selenium IDE. For example, if you use variables for the **type** command, take the following steps:
  - a. Create a **store** command to store a variable in Selenium IDE. Right-click the row where the **type** command is and select **Insert new command**. Click the new row that is inserted and select **store** from the **Command** list. Set the **Value** field with the variable name you want, such as *username*. Leave the **Target** field blank or with a default value.
  - b. Click the row where the **type** command is again and change the **Value** field based on the variable name that you set in the last step. Enclose the variable name in curly brackets ({} ) and precede it

with a dollar sign. Example:

	Command	Target	Value
1	<i>open</i>	/	
2	<i>set window size</i>	1072x680	
3	<i>store</i>	demo	username
4	<i>type</i>	name=q	\${username}
5	<i>send keys</i>	name=q	\${KEY_ENTER}
6	<i>mouse over</i>	linkText=About Selenium	

In this example, the variable *username* is stored with the default value *demo* and the **type** command uses the *username* variable.

3. Create a Selenium script test. You can change the variable values under the **Script variables** section in Cloud App Management console. For more information, see step 4 and step 10 of the topic [“Creating a Selenium script test”](#) on page 633.

### Updating scripts from earlier Selenium IDE versions

You can use `.html` scripts recorded with version of Selenium IDE prior to version 3.12. You will need to edit the `.html` scripts, and save them in the new `.side` format.

Use these considerations to edit your `.html`:

#### Procedure

- Exception: If you want to interact with the `Select2` element, do not use the **select** command (see <https://github.com/SeleniumHQ/selenium-ide>).

The old script is

```
<td>select</td>
<td>id=country</td>
<td>label=United States</td>
```

It should be changed to

```
<tr>
  <td>runScript</td>
  <td>window.scrollTo(0,810)</td>
  <td></td>
</tr>
<tr>
  <td>click</td>
  <td>id=select2-country-container</td>
  <td></td>
</tr>
<tr>
  <td>click</td>
  <td>xpath=//*[@id='select2-country-results']/li[text() = 'United States']</td>
  <td></td>
</tr>
```

## Event generation

In Cloud App Management, each synthetic test generates up to a total of three events. Your test reports the event with the highest severity until the condition that is causing the alert is resolved.

A separate event is raised for three different situations:

### Events based on failed return code

*What is the default behavior?*

If your web application or URL is disabled due to a client or server error, a synthetic test returns a code 400 or above. Every test checks for the response code by default, to determine whether the test is successful or fails.

Each synthetic test automatically triggers an event if a test return code is 400 or above.

This failure detection behavior is enabled by default; no configuration is required. Notice, on the console **Failure detected** is set to **On** by default.

An event is triggered by a single failure on a single Synthetics PoP. Only a successful test result from the same Synthetics PoP can resolve the event.

*How to change the default behavior?*

To disable failure detection, set **Failure detected** to **Off**.

To define the number of consecutive failures that are required before an event is triggered, select **consecutively** under **Failure detected** on the console. The behavior changes so that two consecutive code 400s must be returned before an event is triggered. A synthetic test event is triggered by two consecutive failures on any Synthetics PoP. A successful test from any Synthetics PoP can resolve the event. You can increase the number of consecutive failures that are required to trigger an event.

### Events based on response time

*What is the default behavior?*

Every synthetic test has two built-in conditions that measure response time. If the conditions are not met, a warning or critical event is triggered.

The built-in conditions are:

- Response time >5 seconds triggers a warning event
- Response time >10 seconds triggers a critical event

Events trigger when the synthetic test response time exceeds the default threshold values that are specified in the conditions.

The higher severity overrides the lower one. For example, if the response time is slower than the critical threshold, it is already slower than the warning threshold. Under this circumstance, you see only one event with critical severity. If the synthetic test return code is 400 or above, its response time is ignored.

An event is triggered by one slow response time. A single slow response time event is Synthetics PoP specific. For example, only the same test from the same Synthetics PoP can resolve the event.

*How to change the default behavior?*

To define the number of consecutive slow response times before an event is triggered, select **consecutively** under **Threshold breached**. Then events will be triggered only after two slow response times from any Synthetics PoP. Events can be resolved by a fast response time from any Synthetics PoP. You can increase the number of consecutive slow responses that are needed to trigger an event.

### Events based on content verification

*What is the default behavior?*

You can create customized conditions based on response content validation. You can validate the response from the header response code, header property, or body JSON values.

Select **header response code** to test for one or for a range of HTTP response codes.

Select **header property** to test for a particular HTTP header field property and value.

Select **body json** to test for a particular property from a JSON body.

Events are triggered when the synthetic test content verification result is not met.

The higher severity overrides the lower one. For example, if there are multiple content verification failures, the event always reflects the most critical failure. If the synthetic test return code is 400 or above, its content verification result is ignored. The event is Synthetics PoP specific. For example, only the same test from the same Synthetics PoP can resolve the event.

*How to change the default behavior?*

There is no consecutive event for this kind of detection. The behavior cannot be changed.

For more information and example, see [“Customizing response validation” on page 625](#).

## Viewing Synthetic test results

Use the **Synthetic results** tab in the Cloud App Management console to visualize your synthetic API tests.

### Before you begin








The **Synthetic results** tab is available with the IBM Cloud App Management advanced offering.

### About this task

You must first, install a Synthetics PoP, and create a synthetic test before you can view synthetic test data in the dashboard, for more information, see, [“Synthetics PoP” on page 620](#).

### Procedure

Take these steps in the Cloud App Management console to visualize your synthetic results:

1. Click the **Synthetic results** tab. A table of synthetic results is displayed showing a row for each synthetic test that is created, and including the status of the test. A status icon indicates  Critical,  Warning, or  Normal. The status reflects the highest severity for any events that are triggered by this synthetic test. The synthetics list is sorted from highest severity to lowest. You can also sort by Url or Test name.
2. Click in the Filter test box to search for a synthetic test based on the synthetic test name or url. Click the  to search for synthetics tests.
3. To edit or inspect the test, click .
4. Click  >**Configure**, the **Synthetic tests** page opens. Click on a synthetic test page to edit it. For more information, see [“Synthetics PoP” on page 620](#).
5. Click the synthetic test name or click  >**Inspect**, the following charts are displayed:

#### Events Timeline

This time line chart is summarizing at the synthetic test level. Choose a duration, you can choose; -3hrs, -6hrs, -12hrs, -24hrs, -1week, -2weeks, or -1 months.

If one or more of the synthetic test conditions are not met, or if there is a test failure, an event is raised, and a square with a number will display on the timeline at the times the events were raised. Hover over the number, a list of the events is displayed. From the list, click on an event to open it.

Click anywhere on the timeline to synchronize to that time in the Availability or Response Time charts.

#### Filter by locations

Select one or multiple locations to do the filtering on the locations.



#### Summary

Summary information for the synthetic test across all Synthetics PoP locations is displayed including; Status, Average Response Time, Average response size, Percent available.

#### Availability

Use the Availability chart to view a roll-up of the test instance status for the duration and the Synthetics PoP locations selected. Each line represents a Synthetics PoP location.

For each time point, a Normal or Failed icon is displayed

For each time point, a Failed , or Warning  is displayed. Click an icon to go to the test instance breakdown. Test information is only retained for the previous 24hrs, if you click an icon in a time period previous to this, it will not jump to the test instance breakdown as this information is not available.

#### Response Time

The response time area chart plots the response for each synthetic test for the duration and Synthetics PoP locations selected. The time line is synchronized with the Availability chart. If you

want to see a breakdown of the response time into the component parts for Rest API and Web page Synthetic tests, you must first select just one Synthetics PoP location, then click Breakdown. Each of the following is presented as a percentage of the total response time. The following breakdown is given:

- Blocked
- DNS Resolution
- Connecting
- TLS Setup
- Sending
- Waiting
- Receiving

### Test Instance Breakdown

This table provides a row for each test instance for the Web page and selenium script tests. Click the Download icon to download the HAR, and click the Camera icon to view the screen shot if there is playback failure. You can choose to view 10, 50, or 100 instances.

Sort table headings by: Playback result, Location, Response (ms), Errors, and Timestamp

Collapse a test instance to view summary information for that test instance and a gnatt chart detailing the component values of response breakdown.

Click a Gantt time bar to see a detailed breakdown of the URL.

Click the View errors button to show the details.


## Deleting the Synthetics PoP

You might have multiple Synthetics PoP. If a Synthetics PoP is not needed anymore, follow the steps to delete it.





### Procedure

1. Stop the Synthetics PoP by running the command:

```
./stop-pop.sh
```

2. Click a Synthetic test in Cloud App Management console. In the **Locations** section of the **Edit synthetic test** page, find the Synthetics PoP that you want to delete and click  to delete the Synthetics PoP.

Locations ⓘ

Name	Last seen	Status	Testing from	
			<input type="checkbox"/> Off	
Ontario PoP	Today   2:20 PM	online	<input type="checkbox"/> Off	
PoP from TOR lab	Today   2:26 PM	online	<input type="checkbox"/> Off	
PoP on louie4 worker node	Today   2:26 PM	online	<input type="checkbox"/> Off	
Toronto	Today   2:27 PM	online	<input checked="" type="checkbox"/> On	

**Note:** All the Synthetics tests share the **Locations** information. So you can edit any one Synthetics test to delete the Synthetics PoP and the deletion removes the assignments for all the Synthetics tests that are assigned to this location.

3. A window is shown to notify you of the deletion impact. You can complete the deletion by clicking **Delete** or cancel the deletion by clicking **Cancel**.

## Delete location

×

Are you sure about deleting PoP on louie4 worker node? Deletion is permanent and will remove the location from both the back-end and UI.

Cancel

Delete

### Results

The Synthetics PoP is deleted and you cannot see it in the **Locations** section.

## Monitoring Agent for HMC

The Monitoring Agent for HMC provides you with the capability to monitor the Hardware Management Console (HMC). The HMC agent monitors the availability and health of HMC resources such as CPU, memory, storage, and network. It collects the following metrics: HMC, Managed Server(CEC), LPAR, VIOS, CPUPool, VSCSI, FibreChannel, and NPIV and sends these metrics to the Cloud App Management server.

### Installing and configuring the HMC agent

Install the HMC agent to monitor the availability and health of HMC resources.

#### Procedure

1. If you haven't already downloaded the data collectors installation eImage: `appMgtDataCollectors_2019.4.0.2.tar.gz` (part number CC3FMEN) from [IBM Passport Advantage](#).  
For more information, see “Part numbers” on page 71.
2. Change to the directory where the data collectors installation eImage was downloaded.
3. Extract the HMC agent installation image from the data collectors installation eImage, and go to the `app_mgmt_hmc` directory:

```
tar -zxvf appMgtDataCollectors_2019.4.0.2.tar.gz
cd appMgtDataCollectors_2019.4.0.2
tar -zxvf app_mgmt_hmc.tar.gz
cd app_mgmt_hmc
```

4. Run the installation script:

```
./install.sh
```

When prompted, enter the directory where you want to install the HMC agent or accept the default `/opt/ibm/icam/hmcagent` installation directory.

5. Download the agent configuration package:
  - a) Log in to the Cloud App Management console, click the **Get Started** link on the Welcome page, then select **Administration > Integrations**.
  - b) Click the **New an integration** button.
  - c) In the **Standard monitoring agents** section, select the **Data Collectors Configure** button.
  - d) Select **Download file** and specify the directory where you want to save the compressed ConfigPack file.  
The `ibm-cloud-apm-dc-configpack.tar` file is downloaded.
6. Move to the directory where you installed the HMC agent. It is `/opt/ibm/icam/hmcagent` if you choose the default directory as shown in the following command:

```
cd /opt/ibm/icam/hmcagent
```

7. Set up the Cloud App Management server by running the following script:

```
./setup_icam.sh downloaded_path/ibm-cloud-apm-dc-configpack.tar
```

where *downloaded\_path* is the directory that you downloaded the `ibm-cloud-apm-dc-configpack.tar` file.

**Note:** If the Cloud App Management server changes, the Cloud App Management server setup must be completed again.

#### 8. Configure the connection information to the target HMC:

```
./setup_hmc.sh
```

**Important:** If the target HMC changes or the username and password for the target HMC changes, the HMC connection must be completed again. If you need to complete the HMC connection setup again, complete the following steps:

##### a. Run the `./setup_hmc.sh` script and:

- Enter the hostname or IP address of the HMC to monitor.
- Enter the name of the HMC user.
- Enter the password of the HMC user.

#### 9. Start the HMC agent:

```
./start_hmcagent.sh
```

Other useful commands are:

- To stop the HMC agent, run `./start_hmcagent.sh`
- To check the status of the HMC agent, run `./status_hmcagent.sh`
- To check the date the HMC agent was installed and its version, run `./info_hmcagent.sh`

### Results

The HMC agent installed. The installation log file is located in the `/opt/ibm/icam/hmcagent/log` directory if you choose the default directory during installation or your own directory if you specified one during the installation.

## Uninstalling the HMC agent

You can uninstall the HMC agent.

### Procedure

- Uninstall the HMC agent by running the following script from the agent installation directory, which is `/opt/ibm/icam/hmcagent` if you choose the default directory during installation.

```
./uninstall.sh
```

### Results

The HMC agent and its installation directory are removed.



# Chapter 16. Deploying Unified Agent

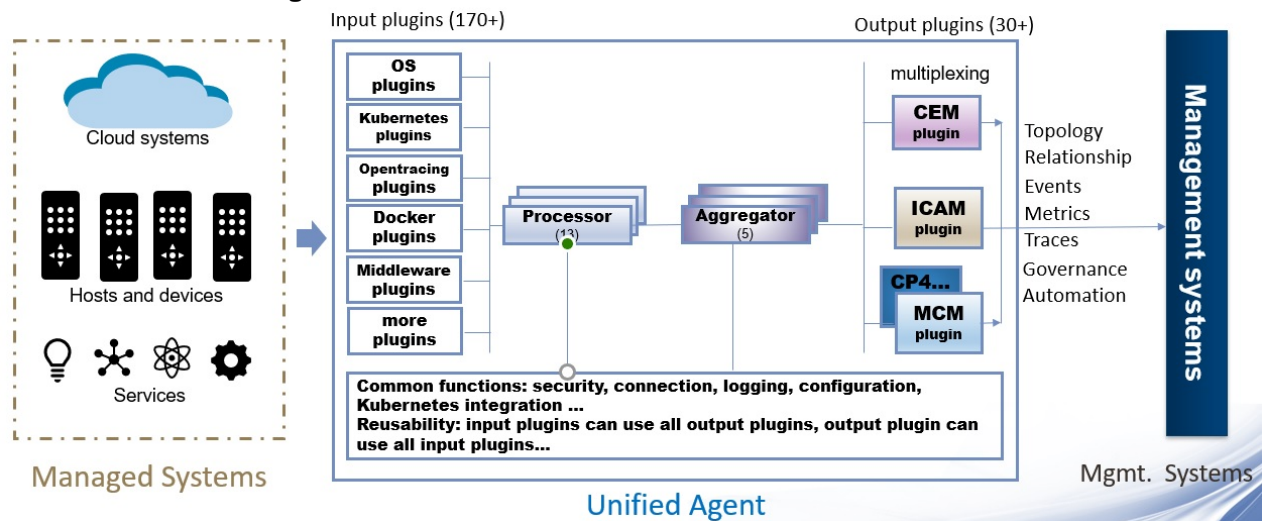
The Unified Agent is an agent for collecting, processing, aggregating, and writing metrics to your IBM Cloud App Management environment. It is based on Telegraf and creates a framework of plug-ins to provide common functions.

## Overview of Unified Agent

The Unified Agent supports receiving OpenTracing workloads including Jaeger and Zipkin, and provides plug-ins to monitor IBM App Connect Enterprise and IBM MQ that are deployed in IBM Cloud Private environment, NGINX and Redis workloads, and IBM API Connect. It also supports plug-ins to monitor OpenShift and DEM.

### Unified Agent Architecture

Unified Agent is based on open source technology called Telegraf. The following picture shows the architecture of Unified Agent.



### Included plug-ins

- Jaeger and Zipkin plug-in
- NGINX plug-in
- Redis plug-in
- IBM API Connect(APIC) plug-in
- IBM App Connect Enterprise(ACE) plug-in
- IBM MQ plug-in
- DEM plug-in
- OpenShift plug-in

## Preparing the deployment of Unified Agent

Before you deploy the Unified Agent, you need to download the Unified Agent Greenfield package and the configuration package, and ensure that your environment meets the prerequisites.

### Before you begin

#### General prerequisites

- Docker 1.7.1 or higher version
- Helm client and server (Tiller) version 2.11.0 or higher

**Note:** In OpenShift environment, you need to manually install Helm client and server, and grant user access to Tiller:

- Do steps 1-3 as stated in the page of [Getting started with Helm on OpenShift](#).
- Grant user access to Tiller by running the following command:

```
oc adm policy add-role-to-user edit username -n namespace
```

Where *username* is the username that you use to deploy the Unified Agent, and *namespace* is the namespace where you install Tiller.

- Make sure that the project where you want to deploy the Unified Agent exists and you have access to this project. Run the following command to grant the access:

```
oc adm policy add-role-to-group edit username -n project
```

Where *project* is the project to deploy the Unified Agent.

- Kubectl client on the environment from where you are installing

#### Special prerequisites for plug-ins

- [UA plug-in for Jagger and Zipkin](#)
- [UA plug-in for NGINX](#)
- [UA plug-in for Redis](#)
- [UA plug-in for IBM API Connect](#)
- [UA plug-in for IBM App Connect Enterprise](#)
- [UA plug-in for IBM MQ](#)
- [UA plug-in for DEM](#)
- [UA plug-in for OpenShift](#)

Table 91. Prerequisites for each plug-in of Unified Agent

Plug-in	System requirements	Other requirements
Jaeger and Zipkin	<a href="#">Link</a>	—
NGINX	<a href="#">Link</a>	Make sure the NGINX monitoring interface is enabled. For more information, see <a href="#">“Enabling the NGINX monitoring interface”</a> on page 648.
Redis	<a href="#">Link</a>	—

<i>Table 91. Prerequisites for each plug-in of Unified Agent (continued)</i>		
<b>Plug-in</b>	<b>System requirements</b>	<b>Other requirements</b>
IBM API Connect	<a href="#">Link</a>	<ul style="list-style-type: none"> <li>The API Connect toolkit is required to provide CLI commands to register the plug-in.</li> <li>The Unified Agent must be deployed in the same Kubernetes environment as IBM API Connect</li> <li>Make sure the service account that you use to install and configure the IBM API Connect plug-in must have access to Kubernetes resources. For more information, see <a href="#">“Authorizing the plug-ins to access Kubernetes resources”</a> on page 651.</li> <li>This plug-in leverages Kubernetes Metrics API to get the pod CPU and Memory usage, so metrics server is required to be deployed in the cluster.</li> </ul>
IBM App Connect Enterprise	<a href="#">Link</a>	<ul style="list-style-type: none"> <li>Make sure the service account that you use to install and configure the IBM App Connect Enterprise plug-in has access to Kubernetes resources. For more information, see <a href="#">“Authorizing the plug-ins to access Kubernetes resources”</a> on page 651.</li> <li>Make sure the Prometheus service for IBM App Connect Enterprise is launched.</li> </ul>
IBM MQ	<a href="#">Link</a>	<ul style="list-style-type: none"> <li>Make sure the service account that you use to install and configure the IBM MQ plug-in has access to Kubernetes resources. For more information, see <a href="#">“Authorizing the plug-ins to access Kubernetes resources”</a> on page 651.</li> <li>Make sure the Prometheus service for IBM MQ is launched.</li> </ul>
DEM	<a href="#">Link</a>	—
OpenShift	<a href="#">Link</a>	—

### About this task

Before proceeding to deploy the Unified Agent, download the Unified Agent greenfield eImage, log in to the Cloud App Management console and download the agent configuration package.

The eImage is the Unified Agent Greenfield package and contains all the installable plug-ins. The configuration package (ConfigPack) contains the configuration files with authentication information required to communicate with the Cloud App Management server and, if the Cloud App Management server is HTTPS enabled, the required certificates.


### Procedure

1. Download the Unified Agent package `unifiedAgent_2019.4.0.1.tar.gz` from [IBM Passport Advantage](#).

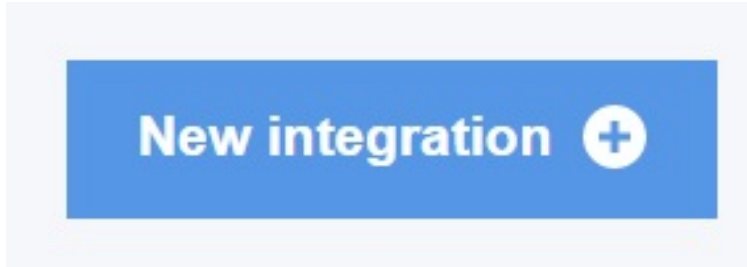
For more information, see [“Part numbers”](#) on page 71.

2. Download the configuration package:

- a) Log in to the Cloud App Management console, click the **Get Started** link on the Welcome page, then select **Administration > Integrations**.



b) Click the **New an integration** button.



c) In the **Standard monitoring agents** section, select the **Data Collectors Configure** button.

d) Select **Download file** and specify the directory where you want to save the compressed ConfigPack file.

The `ibm-cloud-apm-dc-configpack.tar` file is downloaded.

### What to do next

After you complete the preparation steps, you are ready to [deploy the Unified Agent](#).

## Enabling the NGINX monitoring interface

If you want to use the NGINX monitoring by deploying the Unified Agent, you must confirm that the NGINX monitoring interface is enabled.

### Procedure

- Run the following command on the machine where you want to deploy the Unified Agent:

```
http://pod_ip_or_nginx_host:18080/nginx_status
```

where `pod_ip_or_nginx_host` is the fully qualified host name of the NGINX server.

If no status is returned, the NGINX monitoring interface is not enabled.

The NGINX monitoring interface requires loading the `ngx_http_stub_status_module` module. This module helps in collecting basic performance metrics. IBM Cloud Private provides the NGINX Docker image `ibmcom/nginx-ingress-controller`, which has `ngx_http_stub_status` module enabled. If the workloads are using this image, there is no need to do any further configuration. You need only to get the Kubernetes SERVICE or POD IP and verify that the management interface is enabled. For example, `http://NGINX_service_or_node_or_pod_ip:18080/nginx_status`. Some command examples are shown below to help you determine if NGINX workloads are running and to confirm that their management interfaces are enabled.

**Note:** Log into IBM Cloud Private to continue with the following section:

```
> cloudctl login -a <cluster> -u <username>
> kubectl get po -n kube-system -o wide |grep nginx

The command returns output similar to the following

nginx-ingress-controller-jx8vb      1/1      Running    0      5h
10.1.253.201    9.42.75.39

Verify the NGINX management interface status

> curl http://10.1.253.201:18080/nginx_status
```

```
Active connections: 9
server accepts handled requests
5372 5372 22532
Reading: 0 Writing: 2 Waiting: 7
```

Because the pod IP address can change, you can optionally create a service that points to the POD to get a static IP address. The following commands help create a service configuration:

```
> kubectl describe po nginx-ingress-controller-jx8vb -n kube-system

Name:          nginx-ingress-controller-jx8vb
Namespace:     kube-system
Priority:       0
PriorityClassName: <none>
Node:          9.42.75.39/9.42.75.39
Start Time:    Fri, 28 Sep 2018 23:46:39 -0400
Labels:        app=nginx-ingress-controller
```

**Note:** Here the selector is `app=nginx-ingress-controller`. It may differ in your IBM Cloud Private environment

Create a service resource file (`nginx-status.yaml`) as shown here:

```
{
  "apiVersion": "v1",
  "kind": "Service",
  "metadata": {
    "name": "nginx-status",
    "namespace": "kube-system",
    "labels": {
      "app": "nginx-status"
    }
  },
  "spec": {
    "ports": [
      {
        "name": "nginx-status",
        "protocol": "TCP",
        "port": 18080,
        "targetPort": 18080
      }
    ],
    "selector": {
      "app": "nginx-ingress-controller"
    },
    "type": "ClusterIP",
    "sessionAffinity": "None"
  }
}
```

**Note:** Open port 18080 for NGINX status access.

Create the Kubernetes service resource using the file that you created above and obtain the service ip

```
> kubectl create -f nginx-status.yaml
> kubectl describe svc nginx-status -n kube-system

Name:          nginx-status
Namespace:     kube-system
Labels:        app=nginx-status
Annotations:   <none>
Selector:      name=nginx-ingress-controller
Type:          ClusterIP
IP:            10.0.0.243
Port:          nginx-status 18080/TCP
TargetPort:    18080/TCP
Endpoints:     9.37.22.210:18080
Session Affinity: None
Events:        <none>
```

```
curl http://10.0.0.243:18080/nginx_status
```

You can give the NGINX management interface URL `http://10.0.0.243:18080/nginx_status` to the Helm Chart configuration. In on-premises installations, this module is not enabled by default. It

must first be built and then enabled with the configuration parameter `--with-http_stub_status_module`. Please see the NGINX documentation for enablement.

- If you want to monitor NGINX in IBM Cloud Private 3.2.0 or 3.2.1, you need to do extra steps to ensure NGINX monitoring can run successfully.

a) Find the NGINX pod on IBM Cloud Private.

```
# kubectl get po -n kube-system|grep nginx
nginx-ingress-controller-ph8t6           1/1      Running
0                18m
```

b) Export `nginx.tpl`.

```
# kubectl cp kube-system/nginx-ingress-controller-ph8t6:template/nginx.tpl nginx.tpl
```

c) Remove the line `deny all` in `nginx.tpl`.

```
location /nginx_status {
    {{ if $all.Cfg.EnableOpentracing }}
    opentracing off;
    {{ end }}

    {{ range $v := $all.NginxStatusIpv4Whitelist }}
    allow {{ $v }};
    {{ end }}
    {{ if $all.IsIPV6Enabled -}}
    {{ range $v := $all.NginxStatusIpv6Whitelist }}
    allow {{ $v }};
    {{ end }}
    {{ end -}}
    ###This line should be removed , or comment out
    deny all;
    ###End
    access_log off;
    stub_status on;
}
```

d) Create nginx template configmap.

```
kubectl create configmap nginx-template -n kube-system --from-file=nginx.tpl
```

e) Modify nginx daemonset to use this configmap.

```
kubectl edit daemonset nginx-ingress-controller -n kube-system
```

Snippet to add to here:

```
spec:
  containers:
  - args:
    ...
    ...
    ### Here begin the config map setting to copy ###
    volumeMounts:
    - mountPath: /etc/nginx/template
      name: nginx-template-volume
      readOnly: true
    volumes:
    - name: nginx-template-volume
      configMap:
        name: nginx-template
        items:
        - key: nginx.tpl
          path: nginx.tpl
    ### End ###
```

f) Do the verification. After finishing the daemonset edit, the Nginx controller pod will be automatically restarted. If not, manually delete it to take effect.

```
# kubectl get po -n kube-system -o wide|grep nginx
nginx-ingress-controller-ph8t6           1/1      Running
0                3m      10.1.1.13.68      9.46.67.224      <none>           <none>
```

The nginx status is exposed on port 80 according to the nginx config. So, curl this url for verification:

```
#curl http://10.1.13.68/nginx_status
Active connections: 48
server accepts handled requests
 729017 729017 896362
Reading: 0 Writing: 14 Waiting: 33
```

## Authorizing the plug-ins to access Kubernetes resources

To monitor applications running in IBM Cloud Private, the service account that you use to configure the IBM API Connect, IBM App Connect Enterprise and IBM MQ plug-ins must have access to Kubernetes resources through Kubernetes API. Otherwise, you must authorize the service account with appropriate access before you configure the plug-ins.

### About this task

The service account that you use to install and configure the Unified Agent must have access to Kubernetes resources. To determine if the Unified Agent has access to resources, you can use this service account to run the following commands on the Kubernetes master node:

```
kubectl auth can-i list nodes --as system:serviceaccount:
namespace:service_account_name
kubectl auth can-i get pods --as system:serviceaccount:namespace:
service_account_name
kubectl auth can-i list services --as system:serviceaccount:namespace:
service_account_name
```

**Remember:** You must change the *namespace* to the namespace of your environment and the *service\_account\_name* to the name of the service account that you use to configure the Unified Agent. By default, the *service\_account\_name* is default.

See the following example:

```
kubectl auth can-i list nodes --as system:serviceaccount:default:default
kubectl auth can-i get pods --as system:serviceaccount:default:default
kubectl auth can-i list services --as system:serviceaccount:default:default
```

The following procedure authorizes the service account using Role-Based Access Control (RBAC) authorization. For other authorization methods, refer to [Kubernetes documentation](#).

### Procedure

1. Bind the service to a **Role** that has access to query Kubernetes resources in the RBAC mode.
  - a) Create a `rolebinding.yaml` file.

The following example binds the `system:serviceaccount:ops-am:default` account to the `admin` ClusterRole.

```
kind: RoleBinding
apiVersion: rbac.authorization.k8s.io/v1beta1
metadata:
  name: get-pods
  namespace: ops-am
subjects:
- kind: User
  name: system:serviceaccount:ops-am:default
  apiGroup: rbac.authorization.k8s.io
roleRef:
  kind: ClusterRole
  name: admin
  apiGroup: rbac.authorization.k8s.io
```

- b) Run the following command to bind the role:

```
kubectl create -f rolebinding.yaml
```

2. Bind the service to a **ClusterRole** that has access to query Kubernetes resources in the RBAC mode.

- a) Create a `clusterrolebinding.yaml` file.

The following example binds the `system:serviceaccount:ops-am:default` account to the `cluster-admin` ClusterRole.

```
kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1beta1
metadata:
  name: list-cluster
subjects:
- kind: User
  name: system:serviceaccount:ops-am:default
  apiGroup: rbac.authorization.k8s.io
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
```

- b) Run the following command:

```
kubectl create -f clusterrolebinding.yaml
```

## Installing and configuring the Unified Agent

You can deploy the Unified Agent by using Helm chart. In the deployment process, you need to do specific configurations for different plug-ins.

### Before you begin

Ensure that you complete the preparation steps as instructed in [“Preparing the deployment of Unified Agent”](#) on page 646.

### Procedure

1. Extract the `unifiedAgent_2019.4.0.1.tar.gz` package to get cloud monitoring media.

```
tar xzf unifiedAgent_2019.4.0.1.tar.gz
```

2. Log in to your Docker registry.

```
docker login -u my_username -p my_password my_clustername:my_clusterport
```

Where

*my\_username* and *my\_password* are the user name and password for the Docker registry

*my\_clustername* is the name of the cluster that you are monitoring

*my\_clusterport* is the port number for the Docker registry

3. Run the scripts to deploy the Unified Agent. You can choose to install plug-ins `opentracing` (Jaeger and Zipkin), `NGINX`, `Redis`, `IBM APIC`, `IBM ACE`, `IBM MQ`, `DEM`, and `OpenShift` by the option `-pi`.

```
USAGE: deploy.sh
[ -pi Plugin name ] # e.g -pi
opentracing:nginx:redis:ibmmq:ibmapic:ibmace:dem:openshiftua
[ -n Namespace to be deployed ] # e.g -n NameSpace , default is 'ua'
[ -r Release name ] # e.g -r ReleaseName, default is 'monitor'
[ -d Docker repository ]# e.g -d yourcluster.icp:8500, default is 'mycluster.icp:8500'
[ -c Location of configpack zip ]
# e.g -c /Configpack-AbsolutePath/ibm-cloud-apm-dc-configpack.tar
[ -tls <TLS enabled> ] # e.g -tls true or false, default is 'true'
```

Example:

```
./deploy.sh -pi opentracing:ibmmq:ibmapic:ibmace:nginx:redis:dem:openshiftua -r monitor
```



```
-c /tmp/ibm-cloud-apm-dc-configpack.tar
```

In the example, all the plug-ins of Unified Agent are deployed.

**Important:**

- If you deploy the Unified Agent in IBM Cloud Private, make sure you log in to IBM Cloud Private before you perform this step.
- You can select multiple plug-ins to install, but you can deploy only once in one cluster. Next time when you run the `deploy.sh` script, the plug-ins that were previously deployed will be removed and redeployed.
- `-tls` is not mandatory. It is required or not depending on whether `ssl/tls` is enabled in your environment.

4. Configure the plug-ins that you select to install in step 3.

- To configure the Jaeger and Zipkin plug-in, see [“Configuring Jaeger and Zipkin in Unified Agent” on page 654](#).
- To configure the NGINX plug-in, see [“Configuring NGINX monitoring in Unified Agent” on page 655](#).
- To configure the Redis plug-in, see [“Configuring Redis monitoring in Unified Agent” on page 656](#).
- To configure the IBM APIC plug-in, see [“Configuring IBM API Connect monitoring in Unified Agent” on page 656](#).
- To configure the IBM ACE plug-in, see [“Configuring IBM App Connect Enterprise monitoring in Unified Agent” on page 656](#).
- To configure the IBM MQ plug-in, see [“Configuring IBM MQ monitoring in Unified Agent” on page 657](#).
- To configure the DEM plug-in, see [“Configuring DEM in Unified Agent” on page 657](#).
- To configure the OpenShift plug-in, see [“Configuring OpenShift monitoring in Unified Agent” on page 659](#).

5. Validate the deployment.

a) Verify whether pods are successfully started by running the following command:

```
kubectl get po -n namespace |grep ua
```

b) If pods fail to start, check the pod status and pod logs by running the following command:

- Check pod status:

```
kubectl describe po pod-name -n namespace
```

- Check pod logs:

```
kubectl logs pod-name -n namespace
```

c) If you have any issues with the plug-ins, check the plug-in log details:

1) Get the primary pod for the plug-in by running the following command:

```
kubectl describe cm plug-in-configmap -n namespace
```

Where `plug-in-configmap` is the name of the plug-in configmap, for example, it is `ualk-nginx` for NGINX.

Example:

```
kubectl describe cm ualk-nginx -n ua
Name:          ualk-nginx
Namespace:    ua
Labels:       <none>
Annotations:  control-plane.alpha.kubernetes.io/leader:
```

```

        {"holderIdentity":"ua:monitor-ua-cloud-monitoring-
jqhcr","leaseDurationSeconds":60,
"acquireTime":"2019-09-25T03:10:21Z","renewTime":"2019-...
Data
====
Events: <none>

```

In this example, `monitor-ua-cloud-monitoring-jqhcr` is the primary pod.

- d) Detail logs are located at `var/log/ua.log` in container by default. You can change log location by XXX. Run the following command to open the detail log:

```
kubectl exec -it primary-pod -n ua cat var/log/ua.log
```

Where `primary-pod` is the name of the primary pod of the plug-in, for example, `monitor-ua-cloud-monitoring-jqhcr`.

## Results

The Unified Agent plug-ins are installed and begin sending data to the Cloud App Management server for display in the **Resources** dashboard pages.

## Configuring Jaeger and Zipkin in Unified Agent

There are two OpenTracing input plug-ins supporting Zipkin and Jaeger protocols for any custom user apps to connect to OpenTracing service provided by IBM Cloud App Management. When you deploy the Unified Agent, you need to specify some configuration parameter for Jaeger and Zipkin.

### About this task

The Unified Agent as a Daemonset is installed in a namespace named **UA**, and there is a service named **trace**. For all Kubernetes applications, Unified Agent endpoints for OpenTracing can be found in the following table:

Table 92. Unified Agent endpoints for Open tracing	
Protocols	Endpoints
Jaeger	<code>http://trace.ua:14268/api/traces</code>
Zipkin v1	<code>http://trace.ua:9411/api/v1/spans</code>
Zipkin v2	<code>http://trace.ua:9411/api/v2/spans</code>

**Note:** If the user application is not a Kubernetes app, replace `trace.ua` with the host name of Unified Agent.

The Unified Agent can also be used to connect existing tracing systems, for example, Istio trace to Opentracing.

To deploy the Unified Agent with Opentracing plug-ins, do the following steps:

### Procedure

1. Ensure you complete steps 1-3 as instructed in [“Installing and configuring the Unified Agent”](#) on page 652 to deploy the Unified Agent.
2. Specify the Zipkin listening port, default is 9411.
3. Specify the Jaeger listening port, default is 14268.

## Results

Jaeger and Zipkin is successfully installed and configured in Unified Agent.

## What to do next

For normal user instrumented apps with Zipkin or Jaeger protocols, there are no special considerations to use Unified Agent plug-ins for OpenTracing except for specifying the target endpoint. For example, for most Jaeger users, the only task is to set environment variable JAEGER\_ENDPOINT to define the Unified Agent endpoint for Jaeger and let all user code as is.

```
export JAEGER_ENDPOINT=http://trace.ua:14268/api/traces
```

To connect a Spring Boot application to OpenTracing with Unified Agent, do the following steps:

1. Add the following dependencies into your pom.xml file:

```
<dependency>
  <groupId>io.opentracing.contrib</groupId>
  <artifactId>opentracing-spring-web-autoconfigure</artifactId>
  <version>0.3.2</version>
</dependency>

<dependency>
  <groupId>io.jaegertracing</groupId>
  <artifactId>jaeger-core</artifactId>
  <version>0.34.0</version>
</dependency>
```

2. Add the following functions into the SpringBootApplication class:

```
@Bean
public RestTemplate restTemplate(RestTemplateBuilder restTemplateBuilder) {
    return restTemplateBuilder.build();
}

@Bean
public io.opentracing.Tracer tracer() {
    SamplerConfiguration samplerConfig =
    SamplerConfiguration.fromEnv().withType(ConstSampler.TYPE).withParam(1);
    ReporterConfiguration reporterConfig =
    ReporterConfiguration.fromEnv().withLogSpans(true).withSender(
    Configuration.SenderConfiguration.fromEnv().withAgentHost("9.42.82.80").withAgentPort(null));
    Configuration config = new
    Configuration("MySpring").withSampler(samplerConfig).withReporter(reporterConfig);
    return config.getTracer();
}
```

## Configuring NGINX monitoring in Unified Agent

When you deploy the Unified Agent, you can select to install NGINX plug-in to monitor NGINX workloads in Kubernetes environment.

### About this task

To deploy the Unified Agent with NGINX monitoring, do the following steps:

### Procedure

1. Ensure that you complete steps 1-3 as instructed in [“Installing and configuring the Unified Agent”](#) on [page 652](#) to deploy the Unified Agent.
2. Specify the NGINX service address that you want to monitor, for example, `http://nginx_service_IP:18080/nginx_status`.

### Results

NGINX is successfully installed and configured in Unified Agent.

## Configuring Redis monitoring in Unified Agent

When you deploy the Unified Agent, you can select to install Redis plug-in to monitor Redis workloads in Kubernetes environment.

### About this task

To deploy the Unified Agent with Redis monitoring, do the following steps:

### Procedure

1. Ensure you complete steps 1-3 as instructed in [“Installing and configuring the Unified Agent” on page 652](#) to deploy the Unified Agent.
2. Specify the Redis service address that you want to monitor, for example, `tcp://:redisPassword@redis_service_ip:6379`. Your Redis password can often be found by describing your Redis pod to find the name of your Redis password secret, inspecting the yaml of that secret, and decoding the password inside.

### Results

Redis is successfully installed and configured in Unified Agent.

## Configuring IBM API Connect monitoring in Unified Agent

When you deploy the Unified Agent, you can select to install IBM APIC plug-in to monitor IBM API Connect. It can help to determine the health status of the APIC cluster in Kubernetes by retrieving data from Kubernetes API Server, and also gathers APIC Cloud Information from APIC REST APIs including cloud settings, registered services, and cloud events.

### About this task

To deploy the Unified Agent with APIC monitoring, do the following steps:

### Procedure

1. Ensure that you complete steps 1-3 as instructed in [“Installing and configuring the Unified Agent” on page 652](#) to deploy the Unified Agent.
2. Enter the IBM API Connect server hostname of IBM API Connect Cloud Manager, for example, `ui.apic.server.com`.
3. Enter the user name of IBM API Connect server, the default is `admin`.
4. Enter the password of IBM API Connect server.
5. Enter the namespace in which IBM API Connect cluster is deployed. The default is `apiconnect`.
6. Enter the absolute path of the IBM API Connect toolkit to register this plug-in, for example, `/root/apicagent/toolkit/apic-slim`.

### Results

The IBM API Connect plug-in is successfully installed and configured in Unified Agent.

## Configuring IBM App Connect Enterprise monitoring in Unified Agent

When you deploy the Unified Agent, you can select to install IBM ACE plug-in to monitor IBM App Connect Enterprise in IBM Cloud Private cluster environments. It monitors the status of ACE integration server services. You can view information and performance statistics for integration server, message flow, and message flow node in both tabular and chart forms.

### About this task

To deploy the Unified Agent with ACE monitoring, do the following steps:

## Procedure

1. Ensure that you complete steps 1-3 as instructed in [“Installing and configuring the Unified Agent” on page 652](#) to deploy the Unified Agent.
2. Configure the authorization to log in IBM Cloud Private Cluster:
  - a) Enter the IBM Cloud Private cluster user name, the default is admin.
  - b) Enter the IBM Cloud Private cluster password.

## Results

The IBM App Connect Enterprise plug-in is successfully installed and configured in Unified Agent.

## Configuring DEM in Unified Agent

When you deploy the Unified Agent, you can select to install DEM plug-in for digital experience monitoring.

### About this task

To deploy the Unified Agent with DEM, do the following steps:

## Procedure

1. Ensure that you complete steps 1-3 at [“Installing and configuring the Unified Agent” on page 652](#) to deploy the Unified Agent.
2. Specify DEM config parameters.
  - a) Select the jaeger sampler type:
    - `const`: always sampler or always drop
    - `probabilistic`: sampler by probabilistic
    - `rateLimiting`: sampler by count or second
  - b) Select the jaeger sampler param. The valid values for Param field are as follows:
    - For `const` sampler, it is 1 for always true and 2 for always false.
    - For `probabilistic` sampler, it is a probability between 0.0 and 1.0.
    - For `rateLimiting` sampler, it is the number of spans per second.

## Results

The DEM plug-in is successfully installed and configured in Unified Agent.

### What to do next

To enable DEM on IBM HTTP Server or Apache HTTP Server, you must do extra steps to configure the monitoring. For more information, see [“Installing and configuring the DEM plug-in for HTTP Server” on page 670](#).

## Configuring IBM MQ monitoring in Unified Agent

When you deploy the Unified Agent, you can select to install IBM MQ plug-in to monitor IBM MQ container in IBM Cloud Private cluster environments. It can monitor the system resource utilization, such as CPU, memory, and storage. It can also monitor how many API calls failed, how many messages put in and get out from the container IBM MQ service and so on.

### About this task

To deploy the Unified Agent with IBM MQ container monitoring, do the following steps:

## Procedure

1. Ensure that you complete steps 1-3 at [“Installing and configuring the Unified Agent”](#) on page 652 to deploy the Unified Agent.
2. Configure the authorization to log in IBM Cloud Private Cluster.
  - a) Enter the IBM Cloud Private cluster user name, the default is `admin`.
  - b) Enter the IBM Cloud Private cluster password.
3. Configure IBM MQ Services for monitoring. The following config items can be repeated:
  - a) Enter the monitored IBM MQ Service Name. It is the service name to monitor, and it can include asterisk(\*) for fuzzy matching, for example, \* to match all, `abc*` to match the name that begins with `abc`. The default is `*`.
  - b) Enter the Service Namespace. It is the namespace where the services are deployed.
  - c) Enter the IBM MQ administrative REST user name. It is the user name of IBM MQ administration interface and configured when deploying the IBM MQ container. The default is `admin`.  
Where to find the user name and password
  - d) Enter the IBM MQ administrative REST password. It is the password of IBM MQ administration interface and configured when deploying the IBM MQ container.

### Tip:

#### Where to check IBM MQ administrative REST user name and password?

When deploying IBM MQ container, you must create a `Secret` in the target namespace. This must contain the `admin` user password and optionally the app user password to use for messaging. If you do not know the administrative REST user name and password when you deploy the IBM MQ plug-in in Unified Agent, you can do the following steps to check:

- 1) Run the following command to get the secret:

```
kubectl get secret
```

- 2) Find the IBM MQ container secret, run the following command to open it:

```
kubectl edit secret secret_name
```

You can see the password and username information, similar to the following example:

```
"data": {  
  "password": "YWRtaW4=",  
  "username": "YWRtaW4=",  
}
```

- 3) Run the following command to get the password and username:

```
echo password | base64 -d
```

## Results

The IBM MQ plug-in is successfully installed and configured in Unified Agent.

### What to do next

If you want to view the oldest message number of IBM MQ plug-in, do the following steps to set the `MONQ` attribute to `MEDIUM` for `QMGR`:

1. Enter the IBM MQ container by `kubectl exec -it podname /bin/sh -n namespace`.
2. Execute `dspmq` to get the queue manager name.
3. Run `runmqsc queue_manager_name`.
4. Execute `ALERT QMGR MONQ(MEDIUM)`.
5. End from `mqsc` and exit the container.

## Configuring OpenShift monitoring in Unified Agent

When you deploy the Unified Agent, you can select to install OpenShift plug-in to monitor OpenShift route traffic performance and router performance. The plug-in monitors each route response time, volume and error, and also integrates with Kubernetes data collector to enable exploring the associated services and application data.

### About this task

To deploy the Unified Agent with OpenShift monitoring, do the following steps:

### Procedure

1. Ensure that you complete steps 1-3 at [“Installing and configuring the Unified Agent” on page 652](#) to deploy the Unified Agent.
2. When you configure OpenShift parameters, all OpenShift Routers are listed one by one. Confirm whether to monitor the listed router or not. Type y to continue.

### Results

The OpenShift plug-in is successfully installed and configured in Unified Agent.

## Updating the Unified Agent configuration

You can update the Unified Agent configuration by using the configmap.

### About this task

Use the configmap to update the Unified Agent settings.

The following configmaps are created when you deploy the Unified Agent.

```
# kubectl get configmap -n <namespace> | grep ua
cloud-ua-cloud-monitoring-config          1          16h
cloud-ua-cloud-monitoring-pluginconfig    2          16h
cloud-ua-cloud-monitoring-seeelog        1          16h
ualk-ibmace                               0          15d
ualk-ibmapic                              0          15d
ualk-ibmmq                                0          15d
ualk-icam-leader                          0          13d
ualk-nginx                                0          23d
ualk-redis                                 0          23d
```

- `ua-cloud-monitoring-config` is a general config file for Unified Agent, and there is no need to modify it for most scenarios.
- `cloud-ua-cloud-monitoring-pluginconfig` is the file for plug-in config where you can change plug-in setting.
- `cloud-ua-cloud-monitoring-seeelog` is for Unified Agent logging setting where you can change the log level or log path.

Example of see log configuration:

```
<?xml version="1.0" encoding="UTF-8"?>cloud-ua-cloud-monitoring-pluginconfig
<seeelog minlevel="debug">
  <outputs formatid="main">
    <rollingfile type="size" filename="/var/log/ua.log" maxsize="20000000" maxrolls="20" />
  </outputs>
  <formats>
    <format id="main" format="%Date/%Time [%LEV] %Msg%n" />
  </formats>
</seeelog>
```

The log level identifiers for config files are as follows: "trace", "debug", "info", "warn", "error", and "critical".

- The other configmaps are used for debugging when you meet any issues with the plug-ins.

General steps to change a configmap are as follows:

## Procedure

1. Find the correct configmap.

```
kubectl get configmap -n <namespace> |grep ua
```

2. Edit the configmap on demand.

```
kubectl edit configmap <confimap> -n <namespace>
```

To reconfigure each plug-in in Unified Agent, open the configmap `cloud-ua-cloud-monitoring-pluginconfig` to edit.

- For NGINX plug-in, edit the **[[inputs.nginx]]** section. You can change the values to reconfigure the NGINX plug-in.

Example:

```
[[inputs.nginx]]
  ## An array of Nginx stub_status URI to gather stats.
  ### multiple servers can be set as comma-separated list
  ## e.g:
  ##   urls = ['http://10.0.0.0:80/nginx_status', 'http://10.1.2.3:80/nginx_status']
  urls = ['http://10.0.0.0:80/nginx_status']
  ## Use TLS but skip chain & host verification
  #insecure_skip_verify = false
  ## HTTP response timeout (default: 5s)
  response_timeout = "5s"
```

Where:

- `urls` is a comma-separated list of the NGINX server status URLs that you want to monitor.
- `reponse_timeout` is the HTTP response timeout that you want. The default is 5s.
- For Redis plug-in, edit the **[[inputs.redis]]** section. You can change the values to reconfigure the Redis plug-in.

Example:

```
[[inputs.redis]]
  ## specify servers via a url matching:
  ## [protocol://][:password]@address[:port]
  ## e.g.
  ##   tcp://localhost:6379
  ##   tcp://:password@192.168.99.100
  ### multiple servers can be set as comma-separated list
  ## e.g:
  ##   servers =
  ['tcp://:redisPassw0rd@10.0.0.0:6379', 'tcp://:redisPassw0rd@10.1.2.3:6379']
  servers = ['tcp://:redisPassw0rd@10.0.0.0:6379']
  ## Use TLS but skip chain & host verification
  # insecure_skip_verify = true
```

Where `servers` is a comma-separated list of the Redis server URLs that you want to monitor. Your Redis password can often be found by describing your Redis pod to find the name of your Redis password secret, inspecting the yaml of that secret, and decoding the password inside.

- For IBM API Connect plug-in, edit the **[[inputs.ibmapi]]** section. You can change the values to reconfigure the IBM API Connect plug-in.

Example:

```
[[inputs.ibmapi]]
server = "cm.wlavt.com"
username = "admin"
password = "Wmh10DhqawUhCg=="
namespace = "apiconnect"
toolkit_path = "/root/AVT/testcvt/apic-slim"
```



Where:

- `server` is the IBM API Connect server hostname of IBM API Connect Cloud Manager.
- `username` and `password` is the user name and password of the IBM API Connect server.

**Note:** The password in configmap is encoded in base64. If you have changed the password, you need to run `echo password | base64` to generate the base64 password again.

- `namespace` is the namespace where IBM API Connect cluster is deployed, the default value is `apiconnect`.
- `toolkit_path` is the absolute path of the IBM API Connect toolkit to register this plug-in.
- For IBM App Connect Enterprise plug-in, edit the **[[inputs.ibmace]]** section. You can reset the username and password values.

Example:

```
[[inputs.ibmace]]
  username = "admin"
  password = "YWRtaW4K"
```

**Note:** The password in configmap is encoded in base64. If you have changed the password, you need to run `echo password | base64` to generate the base64 password again.

- For IBM MQ plug-in, edit the **[[inputs.ibmmq]]** section.

Example:

```
[[inputs.ibmmq]]
  username = "admin"
  password = "YWRtaW4K"
  mqservices = ["admin:YWRtaW4K@*.*"]
```

Where:

- `username` and `password` is the IBM MQ administrative REST Username and password. Generally speaking, these values will stay unchanged.

**Note:** The password in configmap is encoded in base64. If you have changed the password, you need to run `echo password | base64` to generate the base64 password again.

- `mqservices` contains the service user name, password, service name, and the namespace where the service is deployed. It can include asterisk(\*) for fuzzy matching, for example, \* to match all, `abc*` to match the name that begins with `abc`. The default is \*. You can add, remove, or modify a service or a namespace by changing the `mqservices` value, for example, `["admin:YWRtaW4K@test1-ibm-mq.default, "admin:YWRtaW4K@test2-ibm-mq.default"]`. It means to monitor default namespaces for service `test1-ibm-mq` and `test2-ibm-mq`.

**Note:** The password in configmap is encoded in base64. If you have changed the password, you need to run `echo password | base64` to generate the base64 password again.

- For DEM plug-in, edit the **[[inputs.dem]]** section.

Example:

```
[[inputs.dem]]
  tenant_id = "$APM_TENANT_ID"
  jaeger_sampler_type = "probabilistic"
  jaeger_sampler_param = 0.01
  processor_count = 8
```

Where:

- `tenant_id` is the tenant id of the cluster. Keep the default value.
- `jaeger_sampler_type` is the type of the sampler. The values can be `const`, `probabilistic`, and `rateLimiting`.

- `jaeger_sampler_param` is a value that is passed to the sampler. Valid values for Param field are as follows:
  - For **const** sampler, it is **0** for always false and **1** for always true.
  - For **probabilistic** sampler, the probability value is between **0** and **1**.
  - For **rateLimiting** sampler, it is the number of spans per second.
- `processor_count` is the work thread count of the DEM plugin.
- For OpenShift plug-in, you can run the deployment script `openshiftua_conf.sh` that you can get from the image package.
  - a. Go to the folder where you extract the `unifiedAgent_2019.4.0.tar.gz` package, and find the `openshiftua_conf.sh` file in the subfolder of `./deployment`.
  - b. Execute the script `openshiftua_conf.sh` in the OpenShift environment to generate configuration file.
  - c. During the execution, enter `y` to the router that you want to monitor.
  - d. After the execution is completed, a subfolder `pluginconfig` including the `openshiftua.conf` file is created in the `deployment` directory, for example, `./deployment/pluginconfig/openshiftua.conf`. You can find the following section in the `pluginconfig/openshiftua.conf` file:

```
[[inputs.openshiftua]]
  ## Arrays of Openshift Router metrics URI and other properties to gather route and
  performance data.

  urls = ["http://admin:*****@172.16.174.96:1936/metrics", "http://
admin:*****@9.30.100.180:1938/metrics"]
  routernames = ["router01", "router02"]
  routercpulimits=["100m", "100m"]
  routermemlimits = ["256Mi", "256Mi"]
  routerpods = ["router-3-n4m7a", "routertest-1-9rv6r"]
  routernamespaces = ["default", "default"]
  # providerID from k8monitor configmap
  providerID= ""
  response_timeout = "5s"
```

- e. Open the configmap file `cloud-ua-cloud-monitoring-pluginconfig`, and replace the whole `[[inputs.openshiftua]]` section with the new content that is generated in step d.

```
$ kubectl edit cm ua-cloud-monitoring-pluginconfig -n namespace
```

3. Find the running Unified Agent pod.

```
kubectl get po -n <namespace> |grep ua
```

4. Delete all Unified Agent running pods to restart:

```
kubectl delete po <ua-pod> -n <namespace>
```

5. Optional: Check that the restarted pod is running.

```
kubectl get po -n <namespace> |grep ua
```

## Uninstalling the Unified Agent

You can uninstall the Unified Agent by using Helm chart.

### Procedure

- Run the following command:

```
helm delete ${release-name} --purge --tls
```

**Note:** -tls is not mandatory. It is required or not depending on whether ssl/tls is enabled in your environment.

**Results**

The Unified Agent is removed from your system.



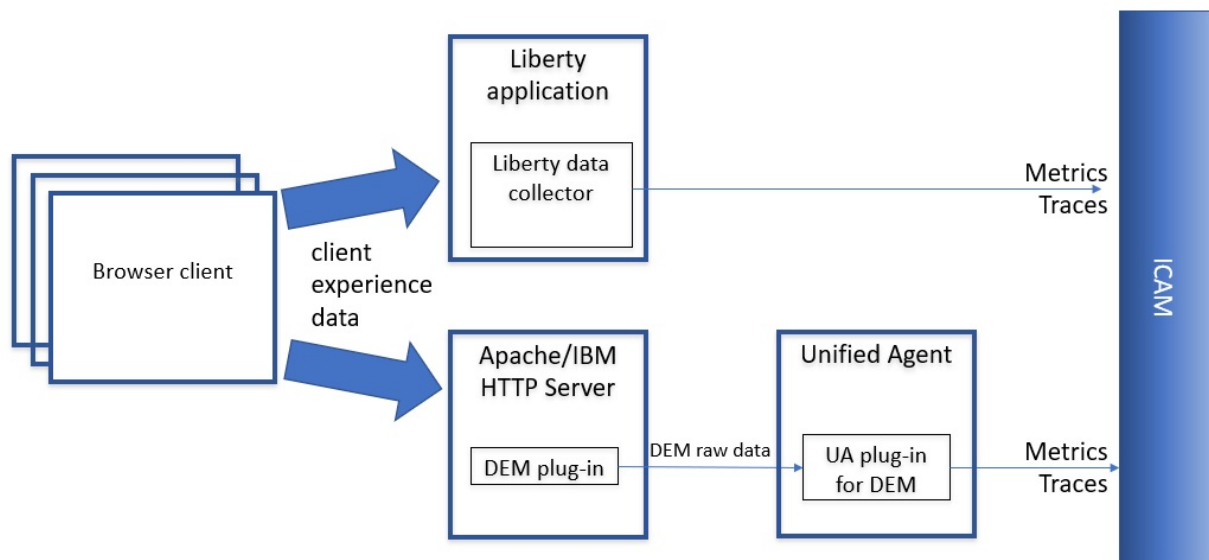
# Chapter 17. Deploying digital experience monitoring(DEM)

Digital experience monitoring(DEM) can be enabled in IBM Cloud App Management to monitor web-based resources and real user experience. It can discover and track traffic, user behavior, and other metrics to help analyze the application performance and usability.

## Overview of DEM

DEM provides a capability to monitor real user experience. It collects data about how actual users interact with and experience web applications. You can enable DEM for Liberty data collector, and you can also enable DEM to monitor HTTP Server.

### Architecture



### Prerequisites

- DEM supports monitoring kube services that are exposed by Ingress only. Make sure that the following settings are enabled in your Kubernetes environment.
  - If SSL Passthrough is enabled in ingress, ensure that X-FORWARDED-FOR is enabled. Set below annotation to ingress to pass client IP to X-FORWARDED-FOR header.

```
ingress.kubernetes.io/configuration-snippet: |
  proxy_set_header X-Forwarded-For $proxy_protocol_addr;
```

- If you are using ingress rewriting rules, ensure that X-Original-URI header is supported to pass original URI. By default it is enabled.
- Make sure the “[Kubernetes data collector](#)” on page 557 is installed and enabled. Otherwise, the DEM service cannot be found in IBM Cloud App Management portal. For more information, see “[Kubernetes data collector](#)” on page 557.

### Deployment options

- To enable DEM in Liberty application monitoring, see “[DEM for Liberty applications](#)” on page 597
- To enable DEM on HTTP Server, see “[Installing and configuring the DEM plug-in for HTTP Server](#)” on page 670.

## DEM for Liberty applications

After you configure the Liberty data collector, you can enable DEM to collect data about how actual users interact with and experience web applications.

### Before you begin

Prerequisites:

- DEM supports monitoring kube services that are exposed by Ingress only. Make sure that the following settings are enabled in your Kubernetes environment.
  - If SSL Passthrough is enabled in ingress, ensure that X-FORWARDED-FOR is enabled. Set below annotation to ingress to pass client IP to X-FORWARDED-FOR header.

```
ingress.kubernetes.io/configuration-snippet: |
  proxy_set_header X-Forwarded-For $proxy_protocol_addr;
```

- If you are using ingress rewriting rules, ensure X-Original-URI header is supported to pass original URI. By default it is enabled.
- Make sure the “Kubernetes data collector” on page 557 is installed and enabled. Otherwise, the DEM service cannot be found in IBM Cloud App Management portal. For more information, see “[Kubernetes data collector](#)” on page 557.
- Make sure the Liberty data collector that you deploy is up-to-date. For more information, see “[Monitoring Liberty applications in Kubernetes environment](#)” on page 591.
- OpenTracing monitoring must be enabled. By default it is enabled in Liberty monitoring.

### About this task

To enable DEM after you configure the Liberty data collector, do the following steps:

### Procedure

1. Open your application deployment yaml file, and add the following environment variables.

```
containers:
  - env:
    - name: IBM_APM_RUM_ENABLED
      value: "true"
```

**Note:** If you want to disable DEM, set the value to `false`.

2. Apply your application deployment yaml file.
3. Optional: If you have configured OpenTracing sampling settings of the Liberty data collector, DEM reads the sampler type and param. If you want to change the sampling settings, open the application deployment yaml file and modify the lines in the **env:** section, for example,

```
- name: JAEGER_SAMPLER_TYPE
  value: "ratelimiting"
- name: JAEGER_SAMPLER_PARAM
  value: "10"
```

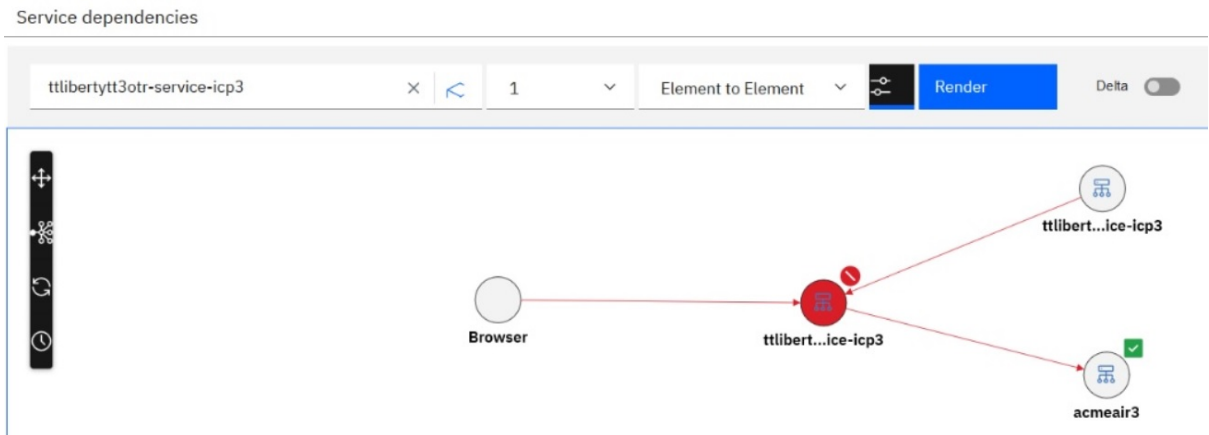
For more information, see “[Customizing the Liberty data collector](#)” on page 597.

### What to do next

Launch a web request, and then do the following steps to verify whether DEM is successfully enabled:

1. In the Cloud App Management console, click the **Resources** tab.
2. Find **Kubernetes Service** from the **All resource types** list and click to open it.
3. Browse the Resource list, click the resource name that you have enabled DEM, and open the resource dashboard.
4. Use one of the following methods to verify:

- Check whether **browser** is displayed in **Service dependencies** section.



- Check whether Browser is listed in the **Related resources** widget.

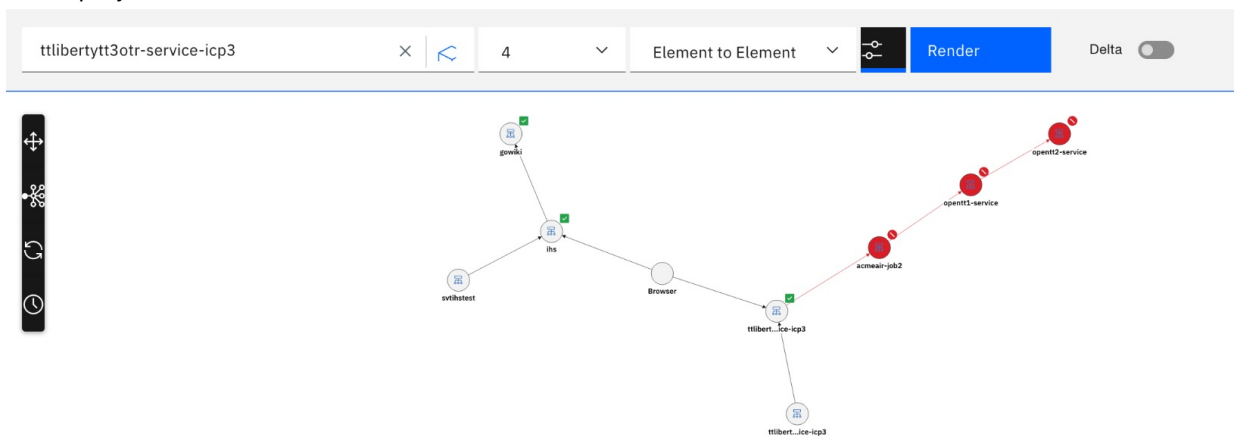
Related resources

Filter

Status	Relation	Resource	Type
✓ Clear	communicatesW	<a href="#">rumdcdev-864768f74b-n67pn (devCluster:default)</a>	Kubernetes Pod
✓ Clear	communicatesW	<a href="#">rumdcdev (devCluster:default)</a>	Kubernetes Ingress
◆ Indeter...	dependsOn/in	<a href="#">Browser</a>	Browser

5. Click to drill down browser to check whether you can get detailed browser data.

**Known limitation:** If multiple services are connected to the Browser in topology, no browser data can be displayed.



## DEM for HTTP server

You can enable DEM for HTTP Server. The DEM plug-in for HTTP Server passively collects data on how actual users are interacting with and experiencing your application. This is achieved through instrumenting the application or injecting code on the page to collect metrics. With DEM and Transaction

Tracking, the DEM plug-in for HTTP Server can monitor web application performance from the browser to the line of code.

### Before you begin

Prerequisites:

- For details about the system requirements, see [SPCR report link](#).
- Software requirements:
  - libstdc++.so.6.0.18 or higher
  - glibc 2.17 or higher
- DEM supports monitoring kube services that are exposed by Ingress only. Make sure that the following settings are enabled in your Kubernetes environment.
  - If SSL Passthrough is enabled in ingress, ensure that X-FORWARDED-FOR is enabled. Set below annotation to ingress to pass client IP to X-FORWARDED-FOR header.

```
ingress.kubernetes.io/configuration-snippet: |
  proxy_set_header X-Forwarded-For $proxy_protocol_addr;
```

- If you are using ingress rewriting rules, ensure X-Original-URI header is supported to pass original URI. By default it is enabled.
- Make sure the “Kubernetes data collector” on page 557 is installed and enabled. Otherwise, the DEM service cannot be found in IBM Cloud App Management portal. For more information, see “Kubernetes data collector” on page 557.

### About this task

To enable DEM for HTTP server, you must first deploy the UA Plug-in for DEM, and then configure the DEM plug-in for HTTP Server. Do the following steps:

### Procedure

1. Deploy the UA Plug-in for DEM. For more information, see “[Installing and configuring the Unified Agent](#)” on page 652.
2. Install and configure the DEM plug-in for HTTP Server. For more information, see “[Installing and configuring the DEM plug-in for HTTP Server](#)” on page 670.

## Installing the UA plug-in for DEM

To enable DEM for HTTP server, you must install the UA plug-in for DEM by deploying the Unified Agent.

### Before you begin

Ensure that you complete the preparation steps as instructed in “[Preparing the deployment of Unified Agent](#)” on page 646.

### Procedure

1. Extract the unifiedAgent\_2019.4.0.1.tar.gz package to get cloud monitoring media.

```
tar xzf unifiedAgent_2019.4.0.1.tar.gz
```

2. Log in to your Docker registry.

```
docker login -u my_username -p my_password my_clustername:my_clusterport
```

Where

*my\_username* and *my\_password* are the user name and password for the Docker registry  
*my\_clustername* is the name of the cluster that you are monitoring  
*my\_clusterport* is the port number for the Docker registry



3. Run the scripts to deploy the Unified Agent. You can choose to install the plug-ins by the option `-pi`.

```
USAGE: deploy.sh
[ -pi Plugin name ] # e.g -pi
opentracing:nginx:redis:ibmmq:ibmapic:ibmace:dem:openshiftua
[ -n Namespace to be deployed ] # e.g -n NameSpace , default is 'ua'
[ -r Release name ] # e.g -r ReleaseName, default is 'monitor'
[ -d Docker repository ]# e.g -d yourcluster.icp:8500, default is 'mycluster.icp:8500'
[ -c Location of configpack zip ]
# e.g -c /Configpack-AbsolutePath/ibm-cloud-apm-dc-configpack.tar
[ -tls <TLS enabled> ] # e.g -tls true or false, default is 'true'
```

Example:

```
./deploy.sh -pi dem -r monitor
-c /tmp/ibm-cloud-apm-dc-configpack.tar
```

In the example, the UA plug-in for DEM are deployed.

### Important:

- If you deploy the Unified Agent in IBM Cloud Private, make sure you log in to IBM Cloud Private before you perform this step.
- You can select multiple plug-ins to install, but you can deploy only once in one cluster. Next time when you run the `deploy.sh` script, the plug-ins that were previously deployed will be removed and redeployed.

4. Configure the DEM plug-in.

a) Select the jaeger sampler type:

- `const`: always sampler or always drop
- `probabilistic`: sampler by probabilistic
- `rateLimiting`: sampler by count or second

b) Select the jaeger sampler param. The valid values for Param field are as follows:

- For `const` sampler, it is 1 for always true and 2 for always false.
- For `probabilistic` sampler, it is a probability between 0.0 and 1.0.
- For `rateLimiting` sampler, it is the number of spans per second.

5. Validate the deployment.

a) Verify whether pods are successfully started by running the following command:

```
kubectl get po -n namespace |grep ua
```

b) If pods fail to start, check the logs for details:

```
kubectl describe po pod-name -n namespace
```

c) If you have any issues with the plug-ins, check the plug-in log details:

1) Get the primary pod for the plug-in by running the following command:

```
kubectl describe cm plug-in-configmap -n namespace
```

Where `plug-in-configmap` is the name of the plug-in configmap.

d) Detail logs are located at `var/log/ua.log` in container by default. You can change log location by XXX. Run the following command to open the detail log:

```
kubectl exec -it primary-pod -n ua cat var/log/ua.log
```

Where `primary-pod` is the name of the primary pod of the plug-in, for example, `monitor-ua-cloud-monitoring-jqhcr`.

## Results

The Unified Agent plug-in for DEM is installed and begins sending data to the Cloud App Management server.

## Installing and configuring the DEM plug-in for HTTP Server

To enable DEM on HTTP Server, you must first deploy the DEM plug-in in Unified Agent, and then install and configure the DEM plug-in for HTTP Server.

### Before you begin

Make sure that the UA plug-in for DEM is deployed in Unified Agent.

### About this task

To deploy the DEM plug-in for HTTP Server, download the plug-in package and then do configurations.

### Procedure

1. Review the part numbers and download `appMgtDataCollectors_2019.4.0.2.tar.gz` from [IBM Passport Advantage](#). For more information, see [Part numbers](#).
2. Extract the package to get the `ibmapm_dem_ihs_datacollector.tgz` package file by running the following command:

```
tar xzf appMgtDataCollectors_2019.4.0.2.tar.gz
cd appMgtDataCollectors_2019.4.0.2
tar xzf app_mgmt_runtime_dc_2019.4.0.2.tar.gz
cd app_mgmt_runtime_dc_2019.4.0.2
```

3. Extract `ibmapm_dem_ihs_datacollector.tgz` and find the files `mod_wrt_2.2.so`, `mod_wrt_2.4.so` and `ihs_dc_enable.sh`.
4. Add the following commands to the Docker file.

- For IBM HTTP Server 8.5.5:

```
COPY mod_wrt_2.2.so /opt/IBM/modules/mod_wrt.so
COPY ihs_dc_enable.sh /opt/IBM/script/ihs_dc_enable.sh
RUN /opt/IBM/script/ihs_dc_enable.sh -c ${HTTPD_CONF_PATH}/httpd.conf
```

- For IBM HTTP Server 9.0:

```
COPY mod_wrt_2.4.so /opt/IBM/modules/mod_wrt.so
COPY ihs_dc_enable.sh /opt/IBM/script/ihs_dc_enable.sh
RUN /opt/IBM/script/ihs_dc_enable.sh -c ${HTTPD_CONF_PATH}/httpd.conf
```

- For Apache HTTP Server 2.2:

```
COPY mod_wrt_2.2.so /opt/IBM/modules/mod_wrt.so
COPY ihs_dc_enable.sh /opt/IBM/script/ihs_dc_enable.sh
RUN /opt/IBM/script/ihs_dc_enable.sh -c ${HTTPD_CONF_PATH}/httpd.conf
```

- For Apache HTTP Server 2.4:

```
COPY mod_wrt_2.4.so /opt/IBM/modules/mod_wrt.so
COPY ihs_dc_enable.sh /opt/IBM/script/ihs_dc_enable.sh
RUN /opt/IBM/script/ihs_dc_enable.sh -c ${HTTPD_CONF_PATH}/httpd.conf
```

- For Oracle HTTP Server 11g:

```
COPY mod_wrt_2.2.so /opt/IBM/modules/mod_wrt.so
COPY ihs_dc_enable.sh /opt/IBM/script/ihs_dc_enable.sh
RUN /opt/IBM/script/ihs_dc_enable.sh -c ${HTTPD_CONF_PATH}/httpd.conf
```

- For Oracle HTTP Server 12c:

```
COPY mod_wrt_2.4.so /opt/IBM/modules/mod_wrt.so
COPY ihs_dc_enable.sh /opt/IBM/script/ihs_dc_enable.sh
RUN /opt/IBM/script/ihs_dc_enable.sh -c ${HTTPD_CONF_PATH}/httpd.conf
```

Where `${HTTPD_CONF_PATH}` is the `httpd.conf` file path.

**Note:**

DEM module depends on `libstdc++.so.6.0.18` or a higher version. Normally `libstdc++.so.6` can be automatically located by `ih_s_dc_enable.sh` during docker build. If the automatic positioning fails, check and install `libstdc++.so.6.0.18` or a higher version.

If the `ih_s_dc_enable.sh` script still fails after `libstdc++.so.6` is installed, specify the `libstdc++.so.6` file path in the Dockerfile, and then run the docker build again. Example:

```
RUN /opt/IBM/script/ih_s_dc_enable.sh -c ${HTTPD_CONF_PATH}/httpd.conf -l /usr/lib/x86_64-linux-gnu/libstdc++.so.6
```

5. Create a ConfigMap file (`env-config.yaml`) like the following:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: env-config
data:
  IBM_APM_DEM_AGENT_ENDPOINT: "uadem.${UA_NAMESPACE}.svc.cluster.local:15001"
  IBM_APM_DEM_ENABLED: "true"
```

Where `${UA_NAMESPACE}` is the namespace where the Unified Agent is installed.

6. Create the Kubernetes ConfigMap resource by using the following command:

```
kubectl apply -f env-config.yaml
```

7. Add the following lines to the deployment yaml file.

```
- volumeMounts:
  - mountPath: /etc/config
    name: config-volume
...
volumes:
- configMap:
  defaultMode: 420
  name: env-config
  name: config-volume
```

For example,

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  annotations:
    deployment.kubernetes.io/revision: "2"
  creationTimestamp: 2018-05-05T13:00:27Z
  generation: 1
  labels:
    app: ih_s-app-selector
    name: ih_s-app-deployment
    namespace: test
spec:
  replicas: 1
  selector:
    matchLabels:
      app: ih_s-app-selector
  strategy:
    rollingUpdate:
      maxSurge: 1
      maxUnavailable: 0
  template:
    metadata:
      creationTimestamp: null
      labels:
        app: ih_s-app-selector
    spec:
      securityContext:
        runAsNonRoot: true
        runAsUser: 1000
        fsGroup: 1000
```

```

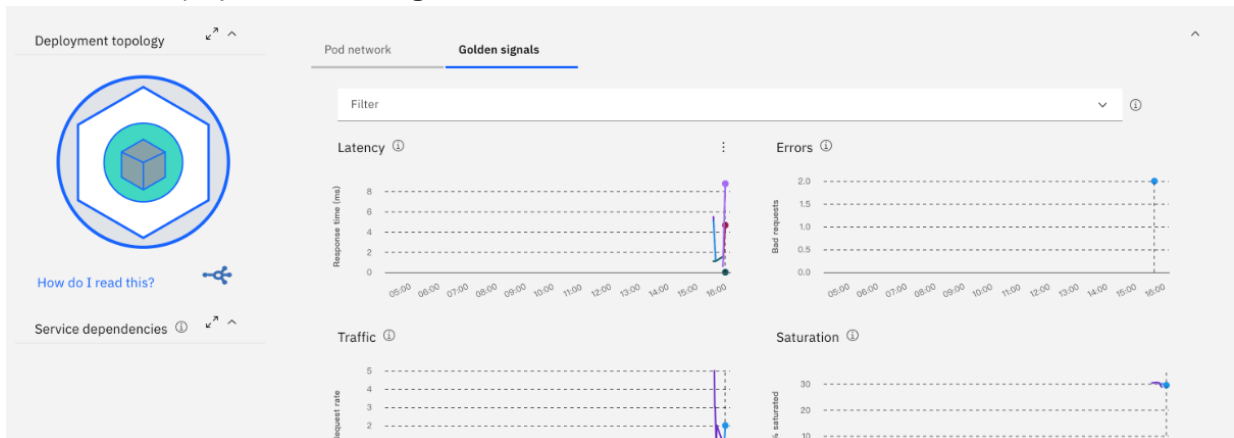
containers:
- volumeMounts:
  - mountPath: /etc/config
    name: config-volume
  image: mycluster.icp:8500/test/ihs:svt
  imagePullPolicy: Always
  name: ihs-app
  resources: {}
  terminationMessagePath: /dev/termination-log
  terminationMessagePolicy: File
volumes:
- configMap:
  defaultMode: 420
  name: env-config
  name: config-volume
  dnsPolicy: ClusterFirst
  restartPolicy: Always
  securityContext: {}
  terminationGracePeriodSeconds: 30

```

## What to do next

Launch a web request, and then do the following steps to verify whether DEM is successfully enabled on HTTP server:

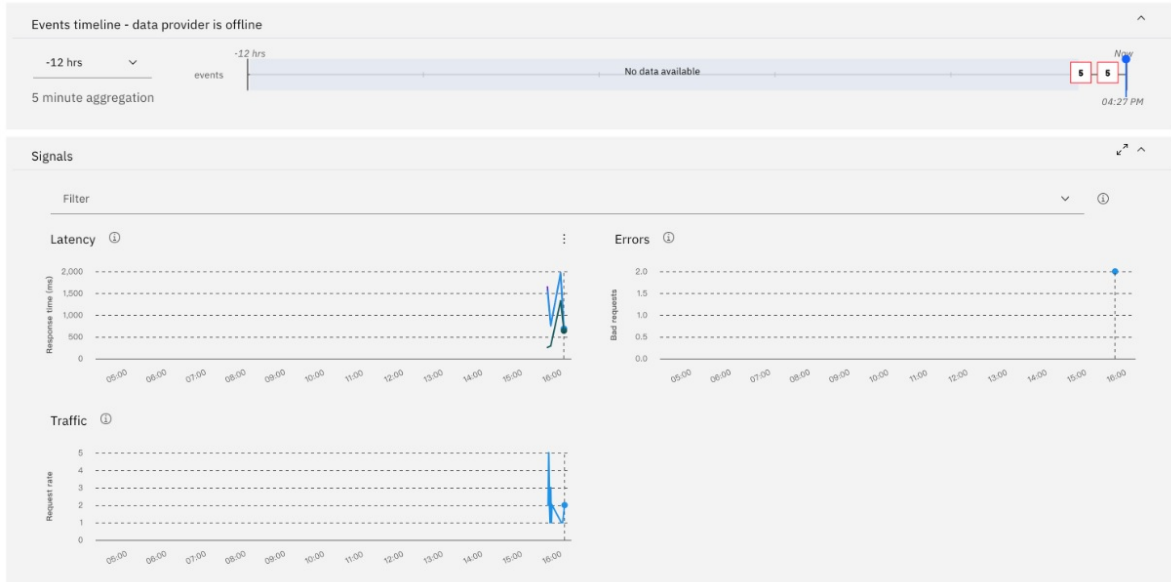
1. In the Cloud App Management console, click the **Resources** tab.
2. Find **Kubernetes Service** from the **All resource types** list and click to open it.
3. Browse the Resource list, click the resource name that you have enabled HTTP Server monitoring, and open the resource dashboard.
4. Metrics are displayed in Golden Signal.



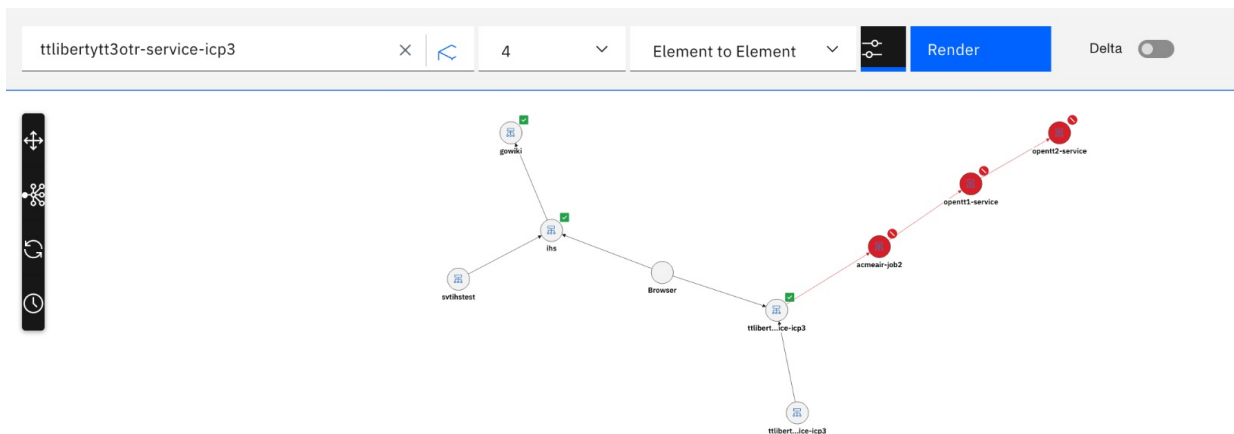
5. Click to drill down browser to check whether you can get detailed browser data.

Resources / Browser

Browser



**Known limitation:** If multiple services are connected to the Browser in topology, no browser data can be displayed.



## Disabling DEM on HTTP server

To disable DEM on HTTP server, you can edit the deployment yaml file.

### About this task

To disable DEM on HTTP server, do the following steps:

### Procedure

1. Open the deployment yaml file and set `IBM_APM_DEM_ENABLED` to `false`:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: env-config
data:
  IBM_APM_DEM_AGENT_ENDPOINT: "uadem.${UA_NAMESPACE}.svc.cluster.local:15001"
  IBM_APM_DEM_ENABLED: "false"
```

2. Update the Kubernetes ConfigMap resource by using the following command:

```
kubect1 apply -f env-config.yaml
```

3. Restart the HTTP Server pods.

## Chapter 18. Integrating with other products

You can integrate other products with Cloud App Management to provide you with a robust monitoring solution.

### Configuring monitoring data sources

Integrate your IBM Tivoli Monitoring and IBM Cloud Application Performance Management supported agents to receive their events and view their resources in the Cloud App Management console. You must define one or more data sources before you configure any other capabilities or see metrics in the **Resources** dashboard page.

### Integrating with IBM Tivoli Monitoring agents

If you have IBM Tivoli Monitoring agents or ITCAM agents (referred to as V6 agents) connecting to Tivoli Enterprise Monitoring Server, you can configure these agents to connect to the Cloud App Management server and then view monitoring data on the Cloud App Management console.

When you configure the V6 agents to connect to the Cloud App Management server, you can specify the following connection modes:

- Use **icam** to connect the V6 agent to Cloud App Management server and disconnect from Tivoli Enterprise Monitoring Server.
- Use **dual** to connect the V6 agent to both Cloud App Management server and Tivoli Enterprise Monitoring Server.
- Use **itm** to connect the V6 agent to Tivoli Enterprise Monitoring Server and disconnect from Cloud App Management server.

See the following table to find out which V6 agents are supported to connect to the Cloud App Management server and which V6 agents support **dual** connection mode.

Agent name	Dual connection mode supported
Cisco UCS agent	No
Citrix VDI agent	No
DataPower agent	Yes
Db2 agent	Yes
HTTP Server agent	Yes
IBM Integration Bus agent	Yes
JBoss agent	Yes
Linux OS agent	Yes
Linux KVM agent	Yes
Microsoft Hyper-V Server agent	No
Microsoft IIS agent	No
Microsoft SQL Server agent	No
NetApp Storage agent	Yes
Oracle Database agent	Yes

Table 93. V6 agents that are supported to connect to the Cloud App Management server (continued)

Agent name	Dual connection mode supported
SAP agent Version 7.1.1 Fix Pack 6 Interim Fix 1	No
SAP HANA Database agent Version 7.1.0 Fix Pack 4	No
UNIX OS agent	Yes
VMware VI agent	Yes
WebSphere Applications agent	Yes
IBM MQ(formerly WebSphere MQ) agent	Yes
Windows OS agent	No

**Known limitations:**

- An agent patch is provided to configure the V6 agent for server connection. However, this agent patch cannot be applied on the system where the Tivoli Enterprise Monitoring Server or the Tivoli Enterprise Portal Server is also installed.
- If you have Tivoli Monitoring private situations, don't try to connect to the Cloud App Management server (don't run the agent2server\_itm script) unless the server has been upgraded to V2019.2.1.1 or later. Otherwise, some configuration files might get deleted.

**Remember:**

- When the supported V6 agent is configured to connect to the Cloud App Management server, unsupported agents that are installed on the same system remain connected to the Tivoli Enterprise Monitoring Server.
- You can reconfigure the V6 agents to connect back to the Tivoli Enterprise Monitoring Server. However, data collected when the V6 agent connects to the Cloud App Management server cannot be retrieved from the Tivoli Enterprise Portal Server.

**Connecting IBM Tivoli Monitoring agents to Cloud App Management server**

To connect IBM Tivoli Monitoring agents to Cloud App Management server, you must first apply an agent patch either locally or remotely to update the agent framework and then configure the agent for server connection.

**Before you begin**

- Download 6.3.0.7-TIV-ITM\_TEMA-IF0008 or a higher agent patch from [IBM Fix Central](#).

- **Linux** **UNIX** (Local configuration only) Determine the architecture of the target operating system to select the appropriate patch file to apply.

**Tip:** Use the `install_dir/bin/cinfo` script to get the architecture code of the operating system.

- (Remote configuration only) Make sure the OS agent is installed on the remote system.
- (WebSphere Applications agent) Transaction tracking data is not yet supported by Cloud App Management. If the V6 agent has been enabled for transaction tracking data collection, reconfigure the V6 agent to disable it before you connect the V6 agent to the Cloud App Management server. For more information, see the V6 agent documentation.
- Be aware of the following limitations before you proceed to apply the agent patch.

**Known limitation:**

- The agent patch cannot be applied on the system where the monitoring server or portal server is also installed.
- After the agent patch is applied, the agent subscription facility (ASF) is started. Many ASF related activities might be logged. You can ignore these messages in logs and no action is required.



## Procedure

- To locally apply the agent patch and configure the agent for server connection, see [“Locally configuring the agent to connect to Cloud App Management server”](#) on page 677.
- To remotely apply the agent patch and configure the agent for server connection, see [“Remotely configuring the agent to connect to Cloud App Management server”](#) on page 679.

### Locally configuring the agent to connect to Cloud App Management server

## Procedure

1. Extract the agent patch 6.3.0.7-TIV-ITM\_TEMA-IF0008.tar or 6.3.0.7-TIV-ITM\_TEMA-IF0008 to the local system where the V6 agent is installed.

In the extracted directory, different fix files are included for all supported operating systems. Use the appropriate file for the target operating system in the following steps.

2. Run the following script to apply the patch.

- **Linux** | **UNIX**

```
cd temp_dir/agent_patch
./install.sh -h install_dir -q -p `pwd`/unix/tfarch.txt
```

- **Windows**

```
cd temp_dir\agent_patch\WINDOWS
setup.exe /w /z"/sf%cd%\deploy\TF_Silent_Install.txt" /s
/f2"install_dir\INSTALLITM\Silent_KTF.log"
```


where:

- *temp\_dir* is the temporary directory that contains the extracted agent patch folder.
- *agent\_patch* is the agent patch file name, for example, 6.3.0.7-TIV-ITM\_TEMA-IF0008.
- *install\_dir* is the V6 agent installation directory. For example, /opt/ibm/itm.
- *arch* is the architecture code of the operating system. Use the appropriate *tfarch.txt* file for the target system, for example, tflx8266.txt.

**Troubleshooting on Windows:** If some product files are locked by other processes on a Windows system, the deployment might fail and the locked files are reported in the Abort IBM Tivoli Monitoring.log file.

To solve this problem, manually stop all processes that are locking the files and try again. For example, if you have WebSphere Applications agent installed, you also need to stop the application server that has the agent data collector installed.

Alternatively, you can add Locked Files=continue to the installation section in the TF\_Silent\_Install.txt and TFX64\_Silent\_Install.txt files within in the *agent\_patch*/WINDOWS/Deploy directory and try again.

For more information about this limitation, see the [Locked files encountered during Windows agent silent installation](#)  technote.

3. Download the agent configuration pack from the Cloud App Management console. The downloaded package contains agent configuration files for server connection.
  - a) Log in to the Cloud App Management console and click **Get Started**.
  - b) Click **Administration > Integrations > Configure an integration (Incoming)**.
  - c) In the Standard monitoring agents section, go to the **ITM/ITCAM Agents** tile and click **Configure**.
  - d) Click **Download file** to download the *ibm-cloud-apm-v6-configpack.tar* file.
4. Extract the .tar file on the systems where the V6 agents are installed.

In the extracted directory, the `.tar` file is for the Linux and UNIX systems and the `.zip` file is used for the Windows systems. Use the appropriate file in the following steps according to the type of the target operating system.

5. If you created private situations, back up your private situation file (`ITM_install_dir/localconfig/pc/pc_situations.xml`).

Otherwise, the private situation file will be lost during reconfiguration of the Tivoli monitoring agent to connect to the Cloud App Management server.

6. To configure the V6 agent for Cloud App Management server connection, extract the `.tar` or `.zip` file, and then run the **agent2server\_itm** script:

- **Linux** | **UNIX**

```
./agent2server_itm.sh -i agent_install_dir -e env.properties -c connection_mode
```

- **Windows**

```
agent2server_itm.bat -i agent_install_dir -e env.properties -c connection_mode
```

where:

- `agent_install_dir` is the V6 agent installation directory, for example, `/opt/ibm/itm`.
- `env.properties` is the path to the file that contains all required server properties. By default, it is the `env.properties` file that is in the same directory of the script file.



**Attention:** This `-e` parameter is optional. You can remove `-e env.properties` from the command if the properties file is in the default directory.

- `connection_mode` is the connection type that you want for the V6 agent. The value can be `icam`, `dual`, or `itm`. If you do not specify any connection type, the default is `icam`.
  - Use `icam` to connect the V6 agent to Cloud App Management server and disconnect from Tivoli Enterprise Monitoring Server.
  - Use `dual` to connect the V6 agent to both Cloud App Management server and Tivoli Enterprise Monitoring Server.
  - Use `itm` to connect the V6 agent to Tivoli Enterprise Monitoring Server and disconnect from Cloud App Management server.

#### Known limitation:

- In `dual` mode, if private situations are defined in IBM Tivoli Monitoring, the private situations run and send events to IBM Tivoli Monitoring EIF destinations. The private situations in IBM Cloud App Management (thresholds run on the agent) will not run.
  - In `icam` mode:
    - If private situations are defined in IBM Tivoli Monitoring, they will not run.
    - The OS agent will still connect to Tivoli Enterprise Monitoring Server to enable remote deploy functions. It will send agent data to Tivoli Enterprise Monitoring Server, and continues to run enterprise situations.
  - In either mode, central configuration files are received from IBM Tivoli Monitoring (if defined) and IBM Cloud App Management in separate folders specific to the server. The files are received no matter whether private situations from that server are run.
7. Optional: If you want to check the current connection mode of the V6 agents, run the following command:

- **Linux** | **UNIX**

```
./agent2server_itm.sh -i agent_install_dir -m
```

- **Windows**

```
agent2server_itm.bat -i agent_install_dir -m
```

**Important:** You can connect the IBM® Tivoli® Monitoring V6 agents to the Cloud App Management server with *icam mode* or *dual mode*.

- If you installed 6.3.0.7-TIV-ITM\_TEMA-IF0008 or earlier, any existing configuration files for Private situations and Central Configuration server that are configured for IBM Tivoli Monitoring are backed up and replaced by those files downloaded from the Cloud App Management server. Therefore, Private situations, Private Historical data, and Central Configuration server files that are configured in IBM Tivoli Monitoring are not available in dual mode.
- If you installed 6.3.0.7-TIV-ITM\_TEMA-IF0009 or later, private situations or private historical data that are configured for IBM Tivoli Monitoring are available in dual mode, but thresholds from the Cloud App Management server are not available. Other files from the Tivoli Enterprise Monitoring Server are processed as usual.

### What to do next

- Log in to the Cloud App Management console to view monitoring data.
- If you are sure that you no longer need to reconnect the agents to Tivoli Enterprise Monitoring Server, remove the offline agents from Tivoli Enterprise Portal.
- If you enable the *icam* connection mode, then, to reconnect the agents to Tivoli Enterprise Monitoring Server, see [“Reconnecting IBM Tivoli Monitoring agents to Tivoli Enterprise Monitoring Server” on page 682](#).
- (IBM MQ(formerly WebSphere MQ) agent only) If you create new agent instances when the agents are connecting to the Cloud App Management server, run the **agent2server\_itm** script again for the new IBM MQ(formerly WebSphere MQ) agent instances to connect to the Cloud App Management server.

### Remotely configuring the agent to connect to Cloud App Management server

#### Procedure

Complete the following steps on a system where the **tacmd** library is available:

1. Extract the agent patch 6.3.0.7-TIV-ITM\_TEMA-IF0008.tar or 6.3.0.7-TIV-ITM\_TEMA-IF0008 to a temporary directory.

There are different .tar files for different operating systems in the extracted agent patch directory. Use the appropriate file for the target operating system in the following steps.

2. On the hub monitoring server system, log in to Tivoli Enterprise Monitoring Server by running the following command from the **tacmd** library:

```
tacmd login -s tems_address -u user_name -p password
```

where:

- *tems\_address* is the host name or IP address of the Tivoli Enterprise Monitoring Server.
- *user\_name* is the user ID that is used to log in to the monitoring server.
- *password* is the user password.

3. Go to the extracted directory that contains the agent patch for the current operating system.

- **Linux** | **UNIX**

```
cd temp/agent_patch/unix
```

- **Windows**

```
cd temp\agent_patch\WINDOWS/Deploy
```

where:

- *temp* is the temporary directory that contains the extracted agent patch folder.
  - *agent\_patch* is the agent patch file name, for example, 6.3.0.7-TIV-ITM-TEMA-IF0008.
4. Run the following command to populate the agent depot:

```
tacmd addbundles -i . -t tf
```

After the command is run, more information about the *tf* component, including its version, is returned.

5. Run the following command from the *tems\_install\_dir/bin* directory to update the agent framework to the version that is returned in Step “4” on page 680.

```
tacmd updateFramework -n node_name -v 063007008
```

where, *node\_name* is the node name of the operating system where the V6 agent is installed.


The following example updates the agent framework on the *kvm-011235:LZ* system.

```
tacmd updateFramework -n kvm-011235:LZ -v 063007008
```

**Troubleshooting on Windows:** If some product files are locked by other processes on a Windows system, the deployment might fail and the locked files are reported in the `Abort IBM Tivoli Monitoring.log` file.

To solve this problem, manually stop all processes that are locking the files and try again. For example, if you have WebSphere Applications agent installed, you also need to stop the application server that has the agent data collector installed.

Alternatively, you can add `Locked Files=continue` to the installation section in the `TF_Silent_Install.txt` and `TFX64_Silent_Install.txt` files within in the *agent\_patch/WINDOWS/Deploy* directory and try again.

For more information about this limitation, see the [Locked files encountered during Windows agent silent installation](#)  technote.

6. Download the agent configuration pack from the Cloud App Management console. The downloaded package contains agent configuration files for server connection.
- a) Log in to the Cloud App Management console and click **Get Started**.
  - b) Click **Administration > Integrations > Configure an integration (Incoming)**.
  - c) In the Standard monitoring agents section, go to the **ITM/ITCAM Agents** tile and click **Configure**.
  - d) Click **Download file** to download the `ibm-cloud-apm-v6-configpack.tar` file.
7. Extract the `.zip` file to the current system.

In the extracted directory, the `.tar` file is for the Linux and UNIX systems and the `.zip` file is used for the Windows system. Use the appropriate file in the following steps according to the type of the target operating system.

**Windows:** If it is difficult to remotely extract the `.zip` file for Windows systems, you can first extract the configuration pack on the current system and then transfer the extracted file to the remote Windows system one by one.

8. On the remote system where the V6 agent is installed, create a temporary directory to save the extracted agent configuration pack. You have multiple ways to do it.

**Example:**

The following example uses the `tacmd executecommand` command to create the `/tmp/configpack` directory on a Linux system as the remote working directory:

```
tacmd executecommand -m kvm-011235:LZ -c "nohup /bin/sh -c 'mkdir /tmp/configpack > /tmp/output'"
```

The following example uses the **tacmd executecommand** command to create the C:\IBM\ITM\configpack directory on a Windows system as the remote working directory:

```
tacmd executecommand -m Primary:IMG-WINDOWS2008:NT -c "md C:\IBM\ITM\configpack"
```

9. Transfer the agent configuration pack to the remote system where the V6 agents are installed.

**Example:**

The following example uses the **tacmd executecommand** command to transfer the linux\_unix\_configpack.tar file from local /mnt/configpacks directory to the /tmp/configpack directory on the remote kvm-011235:LZ system:

```
tacmd putfile -m kvm-011235:LZ -s /mnt/configpacks/linux_unix_configpack.tar -d /tmp/configpack/linux_unix_configpack.tar -t bin
```

The following example uses the **tacmd executecommand** command to transfer the two files extracted from windows\_configpack.zip file from local C:\temp\windows\_configpack directory to the C:\IBM\ITM\configpack directory on the remote Primary:IMG-WINDOWS2008:NT system:

```
tacmd putfile -m Primary:IMG-WINDOWS2008:NT -s C:\temp\windows_configpack\agent2server_itm.bat -d C:\IBM\ITM\configpack\agent2server_itm.bat -t text
tacmd putfile -m Primary:IMG-WINDOWS2008:NT -s C:\temp\windows_configpack\env.properties -d C:\IBM\ITM\configpack\env.properties -t text
```

10. If you created private situations, back up your private situation file (*ITM\_install\_dir/localconfig/pc/pc\_situations.xml*).

Otherwise, the private situation file will be lost during reconfiguration of the Tivoli monitoring agent to connect to the Cloud App Management server.

11. Extract the .tar or .zip file on the remote system if you didn't do it in the previous step, and then run the **agent2server\_itm** script with the -i, -e, and -c options. Use the -i option to specify the agent installation directory, use the -e option to specify the path to the env.properties file in the extracted directory, and use the -c option to specify the connection mode.

**Remember:** On the AIX or Linux system, the sh, bash, or ksh shell is required to run the **agent2server\_itm** script on the remote system.

**Example:**

The following example extracts the /tmp/configpacks/linux\_unix\_configpack.tar file and runs the **agent2server\_itm.sh** script on a Linux system. The V6 agent is installed in the /opt/ibm/itm directory and the sh shell is in the /bin/sh directory. Enable the V6 agent *dual mode* to connect to both Cloud App Management server and Tivoli Enterprise Monitoring Server.

```
tacmd executecommand -m kvm-011235:LZ -c "nohup /bin/sh -c 'sleep 10; tar -xvf /tmp/configpacks/linux_unix_configpack.tar; /tmp/configpacks/agent2server_itm.sh -i /opt/ibm/itm -c dual -e /tmp/configpacks/env.properties > /tmp/output' &" -w /tmp/configpacks
```

The following example directly runs the **agent2server\_itm.bat** script on a Windows system. The V6 agent is installed in the C:\IBM\ITM directory and the sh shell is in the /bin/sh directory. Enable the V6 agent *dual mode* to connect to both Cloud App Management server and Tivoli Enterprise Monitoring Server.

```
tacmd executecommand -m Primary:IMG-WINDOWS2008:NT -c "START /B C:\IBM\ITM\configpack\agent2server_itm.bat -i C:\IBM\ITM -e C:\IBM\ITM\configpack\env.properties -c dual" -w C:\IBM\ITM\configpack
```

## What to do next

- Log in to the Cloud App Management console to view monitoring data.

- If you are sure that you no longer need to reconnect the agents to Tivoli Enterprise Monitoring Server, remove the offline agents from Tivoli Enterprise Portal.
- If you enable the *icam* connection mode, then, to reconnect the agents to Tivoli Enterprise Monitoring Server, see [“Reconnecting IBM Tivoli Monitoring agents to Tivoli Enterprise Monitoring Server”](#) on page 682.
- (IBM MQ(formerly WebSphere MQ) agent only) If you create new agent instances when the agents are connecting to the Cloud App Management server, run the **agent2server\_itm** script again for the new IBM MQ(formerly WebSphere MQ) agent instances to connect to the Cloud App Management server.

### Reconnecting IBM Tivoli Monitoring agents to Tivoli Enterprise Monitoring Server

After the V6 agents are configured to connect to the Cloud App Management server, you can reconfigure them to reconnect to Tivoli Enterprise Monitoring Server again.

#### About this task

Use the **agent2server\_itm** script in the agent configuration packs with *-i* and *-r* options, or with *-i* and *-c* options to reconnect the V6 agents to the Tivoli Enterprise Monitoring Server.

#### Procedure

- On the AIX or Linux system, run one of the following command:

```
./agent2server_itm.sh -i agent_install_dir -r
```

```
./agent2server_itm.sh -i agent_install_dir -c connection_mode
```

Where *connection\_mode* is the connection type that you want for the V6 agent. The value can be *itm* or *dual*.

- Use *itm* to connect the V6 agent to Tivoli Enterprise Monitoring Server and disconnect from Cloud App Management server .
- Use *dual* to connect the V6 agent to both Cloud App Management server and Tivoli Enterprise Monitoring Server.

- On the Windows system, run one of the following command:

```
agent2server_itm.bat -i agent_install_dir -r
```

```
agent2server_itm.bat -i agent_install_dir -c connection_mode
```

Where *connection\_mode* is the connection type that you want for the V6 agent. The value can be *itm* or *dual*.

- Use *itm* to connect the V6 agent to Tivoli Enterprise Monitoring Server and disconnect from Cloud App Management server .
- Use *dual* to connect the V6 agent to both Cloud App Management server and Tivoli Enterprise Monitoring Server.

### Configuring historical data collection for ICAM Agents

You can create history configuration XML files that specify the ICAM Agents to collect data from and send to your Tivoli Data Warehouse. The history file specifies the Warehouse Proxy agent address, the data sets to collect samples from, the frequency of data collection, and how long to keep the data locally. Each history configuration XML file is saved on the Cloud App Management server, which propagates the configuration to all agent instances of this type.

#### Before you begin

Before configuring any resource agent to send data to the Tivoli Data Warehouse, ensure that the equivalent Tivoli Monitoring agent is installed in your Tivoli Monitoring environment. Otherwise, reporting functions can fail.

Install or upgrade your resource agents with the IBM Cloud App Management Version 2019.2.1.1 (or later) agent installation package or apply the IBM Cloud Application Performance Management Version 8.1.4 agent framework interim fix 12 (or later) before configuring historical data collection. Without this update, the historical configurations get lost if the agent is restarted.

### About this task

For every resource agent that can send historical data to Tivoli Data Warehouse, you can create a history configuration file on your Cloud App Management server.

The history configuration file lists the data sets that can send historical data to Tivoli Data Warehouse. If a particular data set that you are interested in does not exist in the sample file, it is likely because this exact data set does not also exist in the Tivoli Monitoring V6.3 agent product or it is not available for historical data collection. You can remove some of the data sets if you do not want to collect data for them.

### Procedure

Create the history configuration XML file from the provided sample:

1. Click one of the following resource agent links and copy the history configuration code block:

[“DataPower agent” on page 686 \(bn\)](#)

[“Db2 agent” on page 686 \(ud\)](#)

[“Hyper-V Server agent” on page 686 \(hv\)](#)

[“IBM Integration Bus agent” on page 687\(qi\)](#)

[“Linux OS agent” on page 687 \(lz\)](#)

[“Microsoft IIS agent” on page 687 \(q7\)](#)

[“Microsoft SQL Server agent” on page 687 \(oq\)](#)

[“UNIX OS agent” on page 688 \(ux\)](#)

[“WebSphere Applications agent” on page 688 \(yn\)](#)

[“IBM MQ\(formerly WebSphere MQ\) agent” on page 688 \(mq\)](#)

[“Windows OS agent” on page 688 \(nt\)](#)

2. Save the code block in a file with the following name:

```
pc_history.xml
```

where *pc* is the two-character agent product code, such as `lz_history.xml` for the Linux OS agent  
Edit the `pc_history.xml` file to configure historical data collection for the resource agent:

3. Specify the Warehouse Proxy agent:

```
<WAREHOUSE LOCATION="ip.pipe:#netaddress[port#]" />
```

where

#### **ip.pipe:**

For non-secure RPC communication between the agent and the Warehouse Proxy Agent, leave at `ip.pipe:.` For secure RPC communication, change to `ip.spipe:.`

#### **#netaddress**

Set the IP address or fully qualified host name of the system where the Warehouse Proxy agent is installed. All hosts where your resource agents run must be able to establish a direct outbound connection to the system using this address or host name.

If you use an IP address, add the `#` sign before the address. If you use a fully qualified host name, make sure the `#` sign is not present before the host name.

#### **port#**

Enter the listening port of the Warehouse Proxy agent. The default port is 63358 for the `ip.pipe` protocol and 65100 for the `ip.spipe` protocol.

You can find the value of the warehouse location string in the RAS1 log file on the Warehouse Proxy agent host. The RAS1 log file is located in the *install\_dir/logs* directory. The file name format is *hostname\_hd\_timestamp-#.log* (for example, *myhost01\_hd\_56d4db3c-01.log*). Search the log file for the `register_interface` message. A RAS1 log message can look like this:

```
"register_interface") Registering "Candle_Warehouse_Proxy": ip.pipe:#9.48.147.34[63358]
```

And the value set in the file can look like this:

```
<WAREHOUSE LOCATION="ip.pipe:#9.48.147.34[63358]" />
```

4. If you want to specify more than one destination or protocol, separate each with a semi-colon (;). For example, you can set the value:

```
<WAREHOUSE LOCATION="ip.spip:#9.11.123.45[65100];ip.pipe:#9.11.123.45[63358];ip.pipe:tdw.example.com[63358]" />
```

In this case, when an agent initiates communications with the Warehouse Proxy agent, it attempts secure RPC communication, then falls back to non-secure RPC communication.

5. Optional: Delete the HISTORY EXPORT rows of the data sets that you do not want to collect history from:

```
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="TABLENAME" />
```

where *TABLENAME* is the data set name.

For example, if you do not want to send Linux\_IP\_Address data samples to the Tivoli Data Warehouse, delete the `<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Linux_IP_Address" />` row.

Data sets are described in the "Attributes" section of the agent help and the reference PDF.

6. In the rows that remain, specify the interval for exporting the data, the interval for collecting the data, and how long to keep the collected samples locally:

#### **EXPORT**

Optional. This parameter specifies the interval in minutes for exporting historical data to the Tivoli Data Warehouse. Valid export intervals are 1, 5, 15, 30, and values divisible by 60; an interval greater than 60 could be 120, 180, 240, and so on, up to 1440. The export interval must also be divisible by the **INTERVAL** parameter value. If you enter an invalid value, no historical data is collected nor exported for the specified attribute group. Default: none.

If used in conjunction with the **USE=A** parameter, the following export integers are valid: 1, 2, 3, 4, 5, 6, 10, 12, and 15.

#### **INTERVAL**

Optional. This parameter specifies the historical data collection interval in minutes. The minimum collection interval is 1 minute and the maximum is 1440 (24 hours). Valid intervals are values that divide evenly into 60 or are divisible by 60: an interval below 60 could be 1, 2, 3, 4, 5, 6, 10, 12, 15, 20, and 30; an interval greater than 60 could be 120, 180, 240, and so on, up to 1440. If you enter an invalid value, no history is collected for the specified attribute group. Default: "15".

#### **RETAIN**

This parameter defines the short-term history data retention period in hours, with a one-hour minimum. There is no limit other than that imposed by storage space on the system. After the retention limit has been reached, the agent deletes oldest data samples as new samples arrive. This retention period ensures that, if the agent loses communication with the Tivoli Data Warehouse for some time, history data is not lost. Default: 6 hours.

7. Save the *pc\_history.xml* file in the Kubernetes master node.
8. Call the Agent Management Services API and enter the following curl command to post the history configuration file to the Cloud App Management server:



```
curl -i -X POST -H "content-type: application/xml" -H "X-TenantID: e69fc647-c775-4131-a48a-7b23455aed78"
-d "@lz_history_config.xml" "http://10.0.0.204:9099/agent_mgmt/0.6/providers/history_configuration?entityType=KLZ"
```

where

*e69fc647-c775-4131-a48a-7b23455aed78* is the cluster tenant ID  
*lz* is the two-character product code for the resource agent  
*10.0.0.204* is the cluster IP address of the agent management service  
*9099* is the agent management service port

You can get the cluster IP address and port number by entering the following command on the IBM Cloud Private master node:

```
kubectl get service | grep agentmgmt
```

9. Repeat these steps for each resource agent that you want to configure for historical data collection.

## Results

After you post the history configuration file to the Cloud App Management server, the server processes the file and distributes the configuration to all online agents of the same type. The time it takes for an agent to receive and process the file and begin historical data collection varies depending on server work load conditions. It might take 15 minutes or more in some cases. As new agents of the applicable type come online, the server automatically distributes the configuration to them.

After your agents receive the configuration, they continue to send history data to the Tivoli Data Warehouse even if connection to the Cloud App Management server is disrupted.

## Example

The following example is the `ud_history.xml` file for the Db2 agent that was configured to collect samples from the KUDINFO00 attributes every 15 minutes, transmits the collected data every hour to the Warehouse Proxy agent at IP address 9.88.765.432, port 63358, and retains the collected data locally for 6 hours:

```
<PRIVATECONFIGURATION>
  <WAREHOUSE LOCATION="ip.pipe:#9.88.765.432[63358]" />
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUDINFO00" />
</PRIVATECONFIGURATION>
```

The `lz_history.xml` historical configuration file that you create from the sample might look like this:

```
<PRIVATECONFIGURATION>
  <WAREHOUSE LOCATION=
"ip.spipe:#9.11.123.45[65100];ip.pipe:#9.11.123.45[63358];ip.pipe:tdw.example.com[63358]" />
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KLZ_CPU" />
  <HISTORY EXPORT="240" INTERVAL="60" RETAIN="24" TABLE="KLZ_DISK" />
  <HISTORY EXPORT="120" INTERVAL="15" RETAIN="6" TABLE="KLZ_SYSTEM_STATISTICS" />
</PRIVATECONFIGURATION>
```

## What to do next

- If you want to update the history configuration for the agent, edit the `pc_history_config.xml` file and post it again to the Cloud App Management server.
- If you want to disable history configuration for an agent type, call the Agent Management Service API and delete the history configuration for that agent. This example uses the Linux OS agent:

```
curl -i -X DELETE -H "X-TenantID: e69fc647-c775-4131-a48a-7b23455aed78"
"http://10.0.0.204:9099/agent_mgmt/0.6/providers/history_configuration?entityType=KLZ"
```

- If you want to view the history configuration for a resource agent type, it is located in the private situation file on the agent machine (example uses Linux OS agent):

```
{Agent_Home}/localconfig/lz/private_situations.xml
```

You can also view the file through the Agent Management Service API with this curl command (example uses Linux OS agent):

```
curl -i -X GET -H "Accept: application/xml" -H "X-TenantID: e69fc647-c775-4131-a48a-7b23455aed78"
"http://10.0.0.204:9099/agent_mgmt/0.6/providers/history_configuration?entityType=KLZ"
```

### Sample history configurations

Use these sample history configuration XML code blocks as a starting point for creating your own history configuration files for each agent type that you want to collect historical data for.

#### DataPower agent

```
<?xml version="1.0" encoding="UTF-8"?><PRIVATECONFIGURATION>
<WAREHOUSE LOCATION="ip.pipe:#netaddress[port#]" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_SYSTEMUPTIME" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_FIRMWAREVERSION" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_DOMAINSTATUS" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_AGENTSTATUS" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_CPUUSAGE" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_MEMORYSTATUS" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_SYSTEMUSAGE" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_SERVICESMEMORYSTATUS" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_TCPTABLE" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_TCPSUMMARY" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_ETHERNETINTERFACE" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_NETWORKTRANSMITDATATHROUGHPUT" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_NETWORKRECEIVEDATATHROUGHPUT" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_HTTPTRANSACTIONS2" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_HTTPMEANTRANSACTIONTIME2" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_MQQUEUEMANAGERS" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_MQCONNECTIONS" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KBN_SQLCONNECTIONS" />
</PRIVATECONFIGURATION>
```

Back to [“Configuring historical data collection for ICAM Agents” on page 682](#)

#### Db2 agent

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<PRIVATECONFIGURATION>
<WAREHOUSE LOCATION="ip.pipe:#netaddress[port#]" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_APPLY_PROGRAM" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_TABLE" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_DCS_DATABASE" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_APPLY_SUBSCRIPTION" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_SYSTEM_RESOURCES" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_IPADDR_TABLE" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_SYSTEM_OVERVIEW" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_APPLICATION00" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_APPLICATION01" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_BUFFER_POOL" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_DATABASE00" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_DATABASE01" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_TABLESPACE" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_DATABASE02" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_LOG" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_LOG_RECORD" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_DIAGNOSTIC_LOG" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_HADR" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_TABLESPACE_AUTO_RESIZE" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KUD_DB2_HADR01" />
</PRIVATECONFIGURATION>
```

#### Hyper-V Server agent

```
<?xml version="1.0" encoding="UTF-8"?>
<PRIVATECONFIGURATION>
<WAREHOUSE LOCATION="ip.pipe:#netaddress[port#]" />
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KHV_HYPER_V_SUMMARY" />
```

```

<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KHV_HYPER_V_SERVER_DISK"/>
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KHV_VIRTUAL_MACHINE"/>
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KHV_VIRTUAL_MACHINE_DETAILS"/>
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KHV_DISK"/>
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KHV_MEMORY"/>
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KHV_PROCESSOR"/>
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6"
TABLE="KHV_HYPER_V_HYPERVISOR_LOGICAL_PROCESSOR"/>
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KHV_VIRTUAL_SWITCH"/>
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KHV_HYPER_V_VIRTUAL_SWITCH"/>
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6"
TABLE="KHV_HYPER_V_VIRTUAL_NETWORK_ADAPTER"/>
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KHV_AVAILABILITY"/>
</PRIVATECONFIGURATION>

```

## IBM Integration Bus agent

```

<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<PRIVATECONFIGURATION>
  <WAREHOUSE LOCATION="ip.pipe:#netaddress[port]"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Broker_Status"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Execution_Group_Status"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Message_Flow_Status"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Accounting_Message_Flow_Statistics"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="JVM_Resource_Statistics"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Accounting_Node_Statistics"/>
</PRIVATECONFIGURATION>

```

## Linux OS agent

```

<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<PRIVATECONFIGURATION>
  <WAREHOUSE LOCATION="ip.pipe:#netaddress[port#]"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KLZ_CPU"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KLZ_DISK"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KLZ_VM_STATS"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KLZ_NETWORK"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KLZ_SYSTEM_STATISTICS"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Linux_IP_Address"/>
</PRIVATECONFIGURATION>

```

## Microsoft IIS agent

```

<?xml version="1.0" encoding="UTF-8"?><PRIVATECONFIGURATION>
  <WAREHOUSE LOCATION="ip.pipe:#netaddress[port#]"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KQ7_ACTIVE_SERVER_PAGES"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KQ7_ASP_NET_APPS_FILTER"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KQ7_IISVRINFO"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KQ7_WPROCESS"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KQ7_WTOTCESS"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KQ7_MEMIISUS"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6"
TABLE="KQ7_INTERNET_INFORMATION_SERVICES_GL"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KQ7_IIS_WEB_SERVER_STATUS"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KQ7_WEB_SERVICE_CACHE"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KQ7_WEB_SERVICE"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KQ7_IIS_WEB_SERVER_SITE_STATUS"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KQ7_FTP_SERVICE"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="KQ7_MICROSOFT_FTP_SERVICE"/>
</PRIVATECONFIGURATION>

```

## Microsoft SQL Server agent

```

<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<PRIVATECONFIGURATION>
  <WAREHOUSE LOCATION="ip.pipe:#netaddress[port#]"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="MS_SQL_SERVER_SUMMARY"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="MS_SQL_DATABASE_SUMMARY"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="MS_SQL_DATABASE_DETAIL"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="MS_SQL_DEVICE_DETAIL"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="MS_SQL_PROCESS_SUMMARY"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="MS_SQL_PROCESS_DETAIL"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="MS_SQL_PROBLEM_SUMMARY"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="MS_SQL_PROBLEM_DETAIL"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="MS_SQL_JOB_SUMMARY"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="MS_SQL_FILEGROUP_DETAIL"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="MS_SQL_LOCK_RESOURCE_TYPE_SUMMARY"/>

```

```
<HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="MS_SQL_DATABASE_MIRRORING"/>
</PRIVATECONFIGURATION>
```

## UNIX OS agent

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<PRIVATECONFIGURATION>
  <WAREHOUSE LOCATION="ip.pipe:#netaddress[port#]"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="SYSTEM"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="DISK"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="UNIX_MEMORY"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="NETWORK"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="SMP_CPU"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="AIX_LPAR"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="AIX_WPAR_INFORMATION"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="AIX_WPAR_CPU"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="AIX_WPAR_PHYSICAL_MEMORY"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="UNIX_IP_Address"/>
</PRIVATECONFIGURATION>
```

## WebSphere Applications agent

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<PRIVATECONFIGURATION REFRESH="Y">
  <WAREHOUSE LOCATION="ip.pipe:#netaddress[port#]"/>
  <HISTORY EXPORT="60" INTERVAL="5" RETAIN="6" TABLE="Application_Server"/>
  <HISTORY EXPORT="60" INTERVAL="5" RETAIN="6" TABLE="Application_Server_Status"/>
  <HISTORY EXPORT="60" INTERVAL="5" RETAIN="6" TABLE="DB_Connection_Pools"/>
  <HISTORY EXPORT="60" INTERVAL="5" RETAIN="6" TABLE="Enterprise_Java_Beans"/>
  <HISTORY EXPORT="60" INTERVAL="1" RETAIN="6" TABLE="Garbage_Collection_Analysis"/>
  <HISTORY EXPORT="60" INTERVAL="5" RETAIN="6" TABLE="Request_Analysis"/>
  <HISTORY EXPORT="60" INTERVAL="5" RETAIN="6" TABLE="Servlets_JSPs"/>
  <HISTORY EXPORT="60" INTERVAL="5" RETAIN="6" TABLE="Thread_Pools"/>
  <HISTORY EXPORT="60" INTERVAL="5" RETAIN="6" TABLE="Web_Applications" />
</PRIVATECONFIGURATION>
```

## IBM MQ(formerly WebSphere MQ) agent

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<PRIVATECONFIGURATION>
  <WAREHOUSE LOCATION="ip.pipe:#netaddress[port]"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Current_Queue_Manager_Status"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Error_Log"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Queue_Data"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Channel_Status"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Queue_Status"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Listener_Status"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Queue_Long-Term_History"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="Channel_Long-Term_History"/>
</PRIVATECONFIGURATION>
```

## Windows OS agent

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<PRIVATECONFIGURATION>
  <WAREHOUSE LOCATION="ip.pipe:#netaddress[port#]"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="NT_LOGICAL_DISK"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="NT_MEMORY_64"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="NT_SYSTEM"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="NT_SERVER"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="NT_PAGING_FILE"/>
  <HISTORY EXPORT="60" INTERVAL="15" RETAIN="6" TABLE="NT_Computer_Information"/>
</PRIVATECONFIGURATION>
```

## Integrating with Cloud APM, Private agents

If you have Cloud APM, Private V8.1.4 agents (referred to as V8 agents) connecting to the on-premises Cloud APM server, you can configure these agents to connect to the Cloud App Management server and then view monitoring data on the Cloud App Management console.

The following V8 agents are supported to connect to the Cloud App Management server.

- Amazon EC2 agent
- Amazon ELB agent

- Azure Compute agent
- Cassandra agent
- Cisco UCS agent
- Citrix VDI agent
- DataPower agent
- DataStage agent\*
- Db2 agent
- HTTP Server agent
- IBM Integration Bus agent
- JBoss agent
- Linux OS agent
- Linux KVM agent
- Microsoft .NET agent
- Microsoft Office 365 agent\*
- Microsoft Hyper-V Server agent
- Microsoft IIS agent
- Microsoft SQL Server agent
- MongoDB agent
- MySQL agent
- NetApp Storage agent
- Oracle Database agent
- PostgreSQL agent
- RabbitMQ agent
- SAP agent
- SAP HANA Database agent
- SAP NetWeaver Java Stack agent
- Skype for Business Server agent
- **2019.4.0.2** Sterling Connect Direct agent
- **2019.4.0.2** Sterling File Gateway agent
- Sybase agent
- Tomcat agent
- VMware VI agent
- UNIX OS agent
- WebLogic agent
- WebSphere Applications agent
- WebSphere Infrastructure Manager agent
- IBM MQ(formerly WebSphere MQ) agent
- Windows OS agent

\*You need to connect at least one agent of this type to the Cloud App Management server before you can connect a v8 agent of this type.

**Remember:**

- When the supported V8 agent is configured to connect to the Cloud App Management server, other unsupported agents installed on the same system are also configured to connect to the Cloud App Management server. However, you cannot view monitoring data from the unsupported agents on the Cloud App Management console.
- You can reconfigure all agents to connect to the Cloud APM server. However, data collected when the V8 agents connect to the Cloud App Management server cannot be retrieved from the Cloud APM console.

### Connecting Cloud APM agents to Cloud App Management server

To connect Cloud APM agents to Cloud App Management server, you must first locally apply an agent patch to update the agent framework and then configure the agent for server connection.

#### Before you begin

- Download the agent patch from [IBM Fix Central](#). Different patches are provided for the following operating systems:
  - AIX: 8.1.4.0-IBM-APM-CORE-FRAMEWORK-AIX-IF0008.tar
  - Linux for System p: 8.1.4.0-IBM-APM-CORE-FRAMEWORK-PLINUXLE-IF0008.tar
  - Linux for System x: 8.1.4.0-IBM-APM-CORE-FRAMEWORK-XLINUX-IF0008.tar
  - Linux for System z: 8.1.4.0-IBM-APM-CORE-FRAMEWORK-ZLINUX-IF0008.tar
  - Windows (32-bit): 8.1.4.0-IBM-APM-CORE-FRAMEWORK-WIN32-IF0008.zip
  - Windows (64-bit): 8.1.4.0-IBM-APM-CORE-FRAMEWORK-WIN64-IF0008.zip
- Diagnostics and transaction tracking data are not yet supported by Cloud App Management. If the V8 agents have been enabled for diagnostics and/or transaction tracking data collection, reconfigure the V8 agents to disable them before you connect the V8 agents to the Cloud App Management server.

#### Procedure

1. Extract the agent patch to the local system where the V8 agent is installed.
2. From the extracted directory, run the following command to apply the patch:

- **Linux** | **UNIX**

```
./apmpatch.sh
```

- **Windows**

```
apmpatch.bat
```

**Note:** If the Cloud APM agents to be connected are not installed in the default directory(/opt/ibm/apm/agent on Linux and AIX, C:\IBM\APM on Windows), you must add the installation path *install\_dir* to the command:

- **Linux** | **UNIX**

```
./apmpatch.sh install_dir
```

- **Windows**

```
apmpatch.bat install_dir
```

The patch will be applied for IBM Monitoring Shared Libraries and IBM GSKit Security Interface.

3. If the V8 agents have been enabled for diagnostics and/or transaction tracking data collection, reconfigure the V8 agents to disable them.

For more information about how to reconfigure the V8 agent, see the [IBM Cloud APM](#) Knowledge Center.

4. Download the agent configuration pack from the Cloud App Management console. The downloaded package contains agent configuration files for server connection.
  - a) Log in to the Cloud App Management console and click **Get Started**.
  - b) Click **Administration > Integrations > New integration**.
  - c) In the Standard monitoring agents section, go to the **APM V8 Agents** tile and click **Configure**.
  - d) Click **Download file** to download the `ibm-cloud-apm-v8-configpack.tar` file.
5. Extract the `.tar` file to the system where the V8 agents are installed.

In the extracted directory, the `.tar` file is for the Linux and UNIX systems, while the `.zip` is for all Windows systems. Use the appropriate file in the following steps according to your operating system.

6. To configure the V8 agent for Cloud App Management server connection, extract the `.tar` and run the **post\_config** script with the `-i` and `-e` options as the user who installed the V8 agent. Use the `-i` option to specify the agent installation directory and use the `-e` option to specify the path to the `env.properties` file in the extracted directory.

- **Linux** | **UNIX**

```
./post_config.sh -i agent_install_dir -e env.properties
```

- **Windows**

```
post_config.bat -i agent_install_dir -e env.properties
```

where `agent_install_dir` is the V8 agent installation directory.

## Results

All V8 agents installed on the same system are configured to connect to the Cloud App Management server. However, you can view monitoring data only for the supported agents on the Cloud App Management console.

## What to do next

- Open the Cloud App Management console to view monitoring data for the supported agents.
- If you are sure that you no longer need to reconnect the agents to the Cloud APM server, remove the offline agents from the Cloud APM console. For more information, see the [Viewing and removing offline agents](#) topic in the IBM Cloud APM Knowledge Center.
- To reconnect the agents to the Cloud APM server, see [“Reconnecting Cloud APM agents to Cloud APM server” on page 691](#).

## Reconnecting Cloud APM agents to Cloud APM server

After the V8 agents are configured to connect to the Cloud App Management server, you can reconfigure them to connect these agents to connect to Cloud APM server again.

## About this task

To reconnect the V8 agents to the Cloud APM server, run the **post\_config** script that is located from the extracted folder of the `ibm-cloud-apm-v8-configpack.tar` file.

You should run these scripts as the user who installed the agent.

## Procedure

- Run the following command:

```
post_config.sh -s cloud_apm_server_address -p cloud_apm_server_protocol -r
```

where:

– `cloud_apm_server_address` is the host name or IP address of the Cloud APM server.

- `cloud_apm_server_protocol` is the protocol of the Cloud APM server. Supported values are `http` and `https`.

### What to do next

After the V8 agents connect to the Cloud APM server, reconfigure the agent or data collector to enable diagnostics and/or transaction tracking again if you still need these data on the Cloud APM console.

## Configuring incoming event sources

The standard integrations are incoming event sources from outside IBM Cloud App Management.

IBM do not provide monitoring agents for the event sources listed below, but do provide the mechanisms to allow the various event sources to forward event data to IBM Cloud App Management via webhooks.

## Creating custom event sources with JSON

You can insert event information into IBM Cloud App Management from any event source that can send the information in JSON format.

### About this task

Using a webhook URL, set your event source to send event information to IBM Cloud App Management. Using an example incoming request in JSON format, define the mapping between the event attributes from your source and the event attributes in IBM Cloud App Management.

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **Webhook** tile and click **Configure**.

4. Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.

**Tip:** Enter a name that identifies the event source you want to receive event information from. A descriptive name will help you identify the event source integration later.

5. Go to your event source and use the generated webhook URL to configure the event source to send event information to IBM Cloud App Management.

**Note:** When IBM Cloud App Management is deployed in an IBM Cloud Private environment, the hostname used in the webhook address (which might be an internal IBM Cloud Private alias) must be resolvable in DNS or the local hosts file where the JSON alerts are being sent from.

6. Copy an incoming JSON request from the event source you are integrating with, and paste it in the **Example incoming request** field of your event source integration in the Cloud Event Management UI.
7. To populate the right event fields in Cloud Event Management from the incoming request, define the mapping between the JSON request attributes and the IBM Cloud App Management event attributes.

**Note:** Four attributes are mandatory as mentioned in this step. You can also set additional attributes to be mapped, as described in the following step.

In the IBM Cloud App Management UI, go to your event source integration and enter values for the event attributes in the **Event attributes** section. Based on this mapping, the IBM Cloud App Management event API then takes values from the incoming request to populate the event information that is inserted into IBM Cloud App Management. For more information about the IBM Cloud App Management API, see **Developer Tools** at <https://console.cloud.ibm.com/apidocs/>.

The following attributes must have a value for an event to be processed by IBM Cloud App Management. Set the mapping in **Event attributes** > **Mandatory event attributes**:



- **Severity:** The event severity level, which indicates how the perceived capability of the managed object has been affected. Values are objects and can be one of the following severity levels: "Critical", "Major", "Minor", "Warning", "Information", "Indeterminate", 60, 50, 40, 30, 20, 10 (60 is the highest and 10 is the lowest severity level).
- **Summary:** String that contains text to describe the event condition.
- **Resource name:** String that identifies the primary resource affected by the event.
- **Event type:** String to help classify the type of event, for example, Utilization, System status, Threshold breach, and other type descriptions.

See later for mapping examples.

**Note:** The event attributes are validated against the mapping to the incoming request example. If the validation is successful, the output is displayed in the **Result** field.

**Important:**

Ensure you are familiar with the JSON format, see <https://www.json.org/>.

For more complex mappings, use JSONATA functions, see <http://docs.jsonata.org/object-functions.html>.

8. **Optional:** In addition to the mandatory attributes, you can set other event attributes to be used, and define the mappings for them. Click **Event attributes > Optional event attributes**, select the additional attributes, and click **Confirm selections**. Then define the mapping between the additional IBM Cloud App Management event attributes and the JSON request attributes to have the correct values populated for the events in IBM Cloud App Management.

**Note:** Most optional attributes can only be added once. Other attributes such as URLs and Related resources can be added more than once. To remove optional attributes, clear the check box for the attribute, or click delete if it has more than one attribute set (for example, URLs), and click **Confirm selections**.

9. Click **Save** to save the event source integration.

**Example**

For examples, see the [Creating custom event sources with JSON](#) topic in the IBM Cloud Event Management Knowledge Center.

**Related information**

## Configuring Amazon Web Services (AWS) as an event source

Amazon Simple Notification Service (SNS) is a web service provided by Amazon Web Services (AWS) that enables applications, end-users, and devices to instantly send and receive notifications from the cloud. You can set up an integration with IBM Cloud App Management to receive notifications from AWS. The Amazon Simple Notification Service (SNS) integration is only available in IBM Cloud App Management, Advanced.


**About this task**

For more information about Amazon SNS, see <https://aws.amazon.com/documentation/sns/>.

Using a webhook URL, alerts generated by AWS monitoring are sent to the IBM Cloud App Management service as events.

**Procedure**

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **Amazon Web Services** tile and click **Configure**.

4. Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.
5. Click **Save**.
6. Log in to your Amazon Web Services account at <https://us-west-2.console.aws.amazon.com/sns/v2/home?region=us-west-2#/topics>
7. Click **Create new topic**, provide a topic name, and click **Create topic**.
8. Go to the **ARN** column in the table and click the link for your topic.
9. Click **Create subscription** and set the fields as follows:
  - a) Select **HTTPS** from the **Protocol** list.
  - b) Paste the webhook URL into the **Endpoint** field. This is the generated URL provided by IBM Cloud App Management.
  - c) Click **Create subscription**.
10. Configure your AWS alarms to send notifications to the Amazon SNS topic you created. The Amazon SNS topic is then used to forward the notification as events to IBM Cloud App Management. For example, you can use Amazon CloudWatch alarms to monitor metrics and send notifications to topics as described in <http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/AlarmThatSendsEmail.html>.
11. To start receiving alert information from AWS, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.


## Configuring AppDynamics as an event source

AppDynamics provides application performance and availability monitoring. You can set up an integration with IBM Cloud App Management to receive alert information from AppDynamics. The AppDynamics integration is only available in IBM Cloud App Management, Advanced.

### About this task

Using a webhook URL, you set up an integration with AppDynamics, and create customized HTTP request templates to post alert information to IBM Cloud App Management based on trigger conditions set in actions as set in AppDynamics policies. The alerts generated by the triggers are sent to the IBM Cloud App Management service as events.

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **AppDynamics** tile and click **Configure**.
4. Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.
5. Click **Save**.
6. Log in to your account at <https://www.appdynamics.com/>.
7. Create a new HTTP request template:
  - a) Click the **Alert & Respond** tab.
  - b) Click **HTTP Request Templates** in the menu bar on the left, and click **New** to add a new template.
  - c) Enter a name for the template.
  - d) In the **Request URL** section, select **POST** from the **Method** list, and paste the webhook URL from IBM Cloud App Management in the **Raw URL** field.

- e) In the **Payload** section, select **application/json** from the **MIME Type** list, and paste the following text in the field:

```
{
  "controllerUrl": "${controllerUrl}",
  "accountId": "${account.id}",
  "accountName": "${account.name}",
  "policy": "${policy.name}",
  "action": "${action.name}",
  #if(${notes})
  "notes": "${notes}",
  #end
  "topSeverity": "${topSeverity}",
  "eventType": "${latestEvent.eventType}",
  "eventId": "${latestEvent.id}",
  "eventGuid": "${latestEvent.guid}",
  "displayName": "${latestEvent.displayName}",
  "eventTime": "${latestEvent.eventTime}",
  "severity": "${latestEvent.severity}",
  "applicationName": "${latestEvent.application.name}",
  "applicationId": "${latestEvent.application.id}",
  "tier": "${latestEvent.tier.name}",
  "node": "${latestEvent.node.name}",
  #if(${latestEvent.db.name})
  "db": "${latestEvent.db.name}",
  #end
  #if(${latestEvent.healthRule.name})
  "healthRule": "${latestEvent.healthRule.name}",
  #end
  #if(${latestEvent.incident.name})
  "incident": "${latestEvent.incident.name}",
  #end
  "affectedEntities": [
  #foreach($entity in ${latestEvent.affectedEntities})
  {
    "entityType": "${entity.entityType}",
    "name": "${entity.name}"
  } #if($foreach.hasNext), #end
  #end
  ],
  "deepLink": "${latestEvent.deepLink}",
  "summaryMessage": "${!latestEvent.summaryMessage.replace('","')}",
  "eventMessage": "${!latestEvent.eventMessage.replace('","')}",
  "healthRuleEvent": ${latestEvent.healthRuleEvent},
  "healthRuleViolationEvent": ${latestEvent.healthRuleViolationEvent},
  "btPerformanceEvent": ${latestEvent.btPerformanceEvent},
  "eventTypeKey": "${latestEvent.eventTypeKey}"
}
```

- f) In the **Response Handling Criteria** section, under **Success Criteria**, click **Add Success Criteria**, and select **200** from the **Status Code** list.
- g) In the **Settings** section, select the **Check One Request Per Event** check box.
- h) Click **Save**.
8. Test your new template. Click **Test**, then click **Add Event Type**, and select an event type. Click **Run Test**. Sample test events are generated and correlated into an incident in IBM Cloud App Management.
- To view the incident and its events, go to the **Incidents** tab on the IBM Cloud App Management UI, click the **All incidents** list, and look for incidents that have a description containing **Cluster: Sample tier**. The event information for these incidents have the event source type set to **AppDynamics**. The event information is available by clicking **Events** on the incident bar, and then clicking the **See more info** button to access all details available for the selected event.
9. Create a new action and add your new template to the action:
- Click **Actions** in the menu bar on the left, and click **Create Action** to add a new template.
  - Select the **Make an HTTP Request** radio button, and click **OK**.
  - Enter a name for the action and select the template you created from the **HTTP Request Template** list.
  - Click **Save**.
10. Add the new action to your AppDynamics policies:

- a) Click **Policies** in the menu bar on the left, and click **Create Policy** to add a new policy, or click **Edit** to edit an existing policy.
  - b) Click **Trigger** in the menu bar on the left, and select the check box for the events that you want to have alerts triggered as part of this policy. The events you select depend on your environment and requirements. For example, you can select all the Health Rule Violation events.
  - c) Click **Actions** in the menu bar on the left, and click **Add**.
  - d) Select **Make an HTTP Request** from the list and click **Select**.
  - e) Click **Save**.
11. To start receiving alert information from the AppDynamics policies based on trigger conditions, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.


## Configuring Datadog as an event source

Datadog provides a monitoring service for your cloud infrastructure. You can set up an integration with IBM Cloud App Management to receive alert information from Datadog. The Datadog integration is only available in IBM Cloud App Management, Advanced.

### About this task

Using a webhook URL, alerts generated by Datadog monitors are sent to the IBM Cloud App Management service as events.

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **Datadog** tile and click **Configure**.
4. Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.
5. Click **Save**.
6. Log in to your account at <http://www.datadoghq.com>.
7. Click **Integrations** in the navigation menu.
8. Go to the **webhooks** tile and click **Install**, or click **Configure** if you already have other webhooks set up.
9. Click the **Configuration** tab, and add a name for the webhook integration in the first available field of the **Name and URL** section at the bottom of the form.
10. Paste the webhook URL into the second field. This is the field after the one where you added the name. This is the generated URL provided by IBM Cloud App Management.
11. Click **Install Integration** or **Update Configuration**, and close the window.
12. Set the webhook for each monitor you want to receive alerts from as follows:
  - a) Click **Monitors > Manage Monitors** in the navigation menu on the left side of the window.
  - b) For existing monitors, hover over the monitor you want to receive alerts from and click **Edit**, or click **New Monitor** if you are setting up a new monitor.
  - c) Go to the **Say what's happening** section and ensure you enter a title for your events in the header text field. For Cluster Alerts, enter a title that includes the following: [Cluster: *resource\_monitored*]. Enter the title in the following format:  
 Title text [Cluster: *resource monitored*]  
 For example: Some of [Cluster: http\_service on redhat] is down.  
 This title is required for the correlation of your Datadog events into incidents.

- d) Go to the main body text field of the **Say what's happening** section, and type @. The available webhook names are listed. Select the name of your webhook integration. The name is also added to the **Notify your team** section.

**Tip:** You can also select your webhook name from the drop-down list in the **Notify your team** section. You can also select users to notify. The selected webhook and users are added to the message in the **Say what's happening** section.

- e) Click **Save**.
  - f) Repeat these steps for each monitor you want to receive alerts from.
13. To start receiving alert information from the Datadog monitors, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.



## Configuring Dynatrace as an event source

Dynatrace provides application performance monitoring. You can set up an integration with IBM Cloud App Management to receive problem notifications from Dynatrace. The Dynatrace integration is only available in IBM Cloud App Management, Advanced.

### About this task

Use a webhook URL and a custom payload to set up the integration between Dynatrace and IBM Cloud App Management.

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **Dynatrace** tile and click **Configure**.
4. Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.
5. Go to step 3. and click  **Copy** to add the custom payload to the clipboard. Ensure you save the custom payload to make it available later in the configuration process. For example, you can save it to a file.
6. Click **Save**.
7. Log in to your account at <https://www.dynatrace.com/> and set up a custom integration:
  - a) Go to **Settings > Integration > Problem notifications**.
  - b) Click **Set up notifications**, and select **Custom integration**.
  - c) On the **Set up custom integration** page, paste the webhook URL from IBM Cloud App Management in the **Webhook URL** field.
  - d) Paste the custom payload from IBM Cloud App Management in the **Custom payload** field.
  - e) Click **Save**.

For more information about setting up custom integrations in Dynatrace, see <https://www.dynatrace.com/support/help/problem-detection/problem-notification/how-can-i-set-up-outgoing-problem-notifications-using-a-webhook/>.

8. Set the alerting rules for Availability, Error, Slowdown, Resource and Custom alerts in Dynatrace as described in <https://www.dynatrace.com/support/help/problem-detection/problem-notification/how-can-i-filter-problem-notifications-with-alerting-profiles/>. The alerting rules determine what problem notifications are sent to IBM Cloud App Management as events.
9. Set the anomaly detection sensitivity for infrastructure components in Dynatrace as described in <https://www.dynatrace.com/support/help/problem-detection/anomaly-detection/how-do-i-adjust->

[anomaly-detection-for-infrastructure-components/](#). The detection sensitivity and alert thresholds determine what problem notifications are sent to IBM Cloud App Management as events.

10. To start receiving problem notifications as events from Dynatrace, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.

## Configuring Elasticsearch as an event source

Elasticsearch is a distributed, RESTful search and analytics engine that stores data as part of the Elastic Stack. You can set up an integration with Elasticsearch to send log information to IBM Cloud App Management as events. The Elasticsearch integration is only available in IBM Cloud App Management, Advanced.


### Before you begin

Ensure you have the X-Pack extension for the Elastic Stack installed as described in <https://www.elastic.co/guide/en/x-pack/current/installing-xpack.html>.

### About this task

Using the X-Pack Alerting (via Watcher) feature, you configure watches to send event information to IBM Cloud App Management. For information about X-Pack Alerting via Watcher, see <https://www.elastic.co/guide/en/x-pack/current/how-watcher-works.html>.

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **Elasticsearch** tile and click **Configure**.
4. Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.
5. Click **Save**.
6. Configure the X-Pack watcher feature in Elasticsearch to forward events to IBM Cloud App Management. For example, to configure the watcher using the Kibana UI:
  - a) Log in to the Kibana UI and access the Watcher UI as described in <https://www.elastic.co/guide/en/kibana/current/watcher-getting-started.html>.

If you are using IBM Cloud Private, you can configure the included Elasticsearch engine to send events to IBM Cloud App Management. You can open the Kibana UI from the navigation menu in IBM Cloud Private by clicking **Network Access** > **Services** > **Kibana**, or by clicking **Platform** > **Logging**.

**Note:** Ensure you have Kibana installed in IBM Cloud Private as described in [https://www.ibm.com/support/knowledgecenter/en/SSBS6K\\_2.1.0.3/featured\\_applications/kibana\\_service.html](https://www.ibm.com/support/knowledgecenter/en/SSBS6K_2.1.0.3/featured_applications/kibana_service.html).

- b) Create a new advanced watch as described in <https://www.elastic.co/guide/en/kibana/current/watcher-create-advanced-watch.html>. Update the fields as follows:
  - Enter an ID and name.
  - Configure your watch definition based on your requirements and add it to the **Watch JSON** field. For more information, see <https://www.elastic.co/guide/en/x-pack/6.2/how-watcher-works.html#watch-definition>.
  - Paste the webhook URL from IBM Cloud App Management in the `url` field under the actions settings.

The following is an example watch definition for IBM Cloud Private environments where the watch is triggered every 5 minutes to load the Logstash logs that were written in the last 5 minutes and

contain any of the following keywords: failed, error, or warning. The watcher posts the payload for such logs to IBM Cloud App Management using the webhook URL.

```
{
  "trigger": {
    "schedule": {
      "interval": "5m"
    }
  },
  "input": {
    "search": {
      "request": {
        "indices": [
          "logstash-2018*"
        ],
        "body": {
          "query": {
            "bool": {
              "must_not": {
                "match": {
                  "kubernetes.container_name": "custom-metrics-adapter"
                }
              }
            }
          }
        },
        "filter": [
          {
            "range": {
              "@timestamp": {
                "gte": "now-5m"
              }
            }
          },
          {
            "terms": {
              "log": [
                "failed",
                "error",
                "warning"
              ]
            }
          }
        ]
      }
    }
  },
  "actions": {
    "my_webhook": {
      "webhook": {
        "method": "POST",
        "headers": {
          "Content-Type": "application/json"
        },
        "url": "<CEM WEBHOOK>",
        "body": "{{#toJson}}ctx.payload{{/toJson}}"
      }
    }
  }
}
```

**Important:** Ensure you set the trigger for the watch to a frequency that suits your requirements for monitoring the logs. Consider the load on the system when setting frequency. In the previous example, the watch is triggered every 5 minutes to load the logs that were written in the last 5 minutes using the "schedule": {"interval": "5m"} and "@timestamp": {"gte": "now-5m"} settings. If you set interval to less than 5 minutes in this case, then the same logs are sent to IBM Cloud App Management more than once, repeating event data in the correlated incidents.

**Restriction:** The "terms": {"log": []} section in the watch definition determines the mapping to the event severity levels in IBM Cloud App Management. The default values are "failed", "error", and "warning", and are mapped to "critical", "major", and "minor" severity levels. If you use any other value, the event severity is mapped to "indeterminate" in IBM Cloud App Management.



**Attention:** In IBM Cloud Private environments ensure you exclude "kubernetes.container\_name": "custom-metrics-adapter" from your watch definition using the following setting:

```
"must_not": {  
  "match": {  
    "kubernetes.container_name": "custom-metrics-adapter"  
  }  
}
```

The size of the custom-metric-adapter logs can be large and overload the Cloud Event Management processing. In addition, the log format is unreadable to users.

- c) Save the watch.
7. If you are using IBM Cloud Private, ensure the X-Pack watcher feature is enabled; for example:
  - a) Load the ELK (Elasticsearch, Logstash, Kibana) stack ConfigMap into a file using the following command:

```
kubectl get configmaps logging-elk-elasticsearch-config --namespace=kube-system -o yaml > elasticsearch-config.yaml
```
  - b) Edit the elasticsearch-config.yaml file to enable the watcher: `xpack.watcher.enabled: true`
  - c) Save the file, and replace the ConfigMap using the following command:

```
kubectl --namespace kube-system replace -f elasticsearch-config.yaml
```
  - d) Restart Elasticsearch and Kibana.
8. To start receiving log information as events from Elasticsearch, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.

## Configuring Jenkins as an event source

Jenkins helps automate software development processes such as builds to allow continuous integration. You can set up an integration with IBM Cloud App Management to receive notifications about jobs from Jenkins projects. The Jenkins integration is only available in IBM Cloud App Management, Advanced.

### Before you begin

If you are using IBM Cloud App Management in an IBM Cloud Private environment, your CA certificate might need to be an X.509 certificate. Complete these steps to convert your PEM certificate:

1. Run the following command:

```
openssl pkcs7 -in cert.pem -out cert.crt -print_certs
```

2. Import your certificate to the JVM keystore as a trusted certificate:

```
keytool -storepass <store_password> -import -noprompt -trustcacerts -alias  
<certificate_alias>  
-keystore cacerts -file cert.crt
```

3. Restart your Jenkins server process to pick up the new certificate.
4. Ensure your Jenkins server host can resolve the domain name of your IBM Cloud App Management installation.
5. Modify the DNS server or add the host and domain name to the hosts file.


### About this task

Notifications can be sent for single job stages or all stages of a job. Configure each project separately from which you want to receive notifications. The notifications are raised in IBM Cloud App Management as events. The events are then correlated into incidents.

**Important:** The Jenkins server needs the Notification Plug-in to send the notifications.



## Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **Jenkins** tile and click **Configure**.
4. Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.
5. Click **Save**.
6. Log into your Jenkins server as administrator.
7. Ensure that the Notification Plug-in is installed on your Jenkins server.

**Tip:** Check first whether the plug-in is installed by clicking **Jenkins > Manage Jenkins > Manage Plugins**. Go to the **Installed** tab and look for the **Notification plugin**. If not in the list of installed plug-ins, go to the **Available** tab and search for **Notification plugin**. Select the check box for the plug-in and click **Install**.

8. Configure the Jenkins project you want to receive notifications from as follows:
  - a) Click the project name and then click **Configure**.
  - b) Click the **Job Notifications** tab, and click **Add Endpoint**.
  - c) Set up the connection as follows:
    - Select **JSON** from the **Format** list.
    - Select **HTTP** from the **Protocol** list.
    - Select when you want to receive notifications about the job from the **Event** list. For example, **All Events** sends a notification for each job phase, while **Job Finalized** only triggers a notification when the job has completed, including post-build activities. Select **All Events** to receive detailed information about the jobs.
    - Paste the webhook URL into the **URL** field. This is the generated URL provided by IBM Cloud App Management.
    - Enter **5** in the **Log** field. This determines the number of lines to include from the log in the message.
  - d) Click **Save**
  - e) Repeat the steps for each project you want to receive notification from.
9. To start receiving notifications about Jenkins jobs, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.

## Configuring Logstash as an event source

You can forward log data to IBM Cloud App Management from Logstash. The Logstash integration is only available in IBM Cloud App Management, Advanced.

### Before you begin


By default, the IBM Cloud Private installer deploys an Elasticsearch, Logstash and Kibana (ELK) stack to collect system logs for the IBM Cloud Private managed services, including Kubernetes and Docker. For more information, see [https://www.ibm.com/support/knowledgecenter/en/SSBS6K\\_2.1.0.2/manage\\_metrics/logging\\_elk.html](https://www.ibm.com/support/knowledgecenter/en/SSBS6K_2.1.0.2/manage_metrics/logging_elk.html)

**Note:** Ensure you meet the prerequisites for IBM Cloud Private, such as installing and configuring the kubectl, the Kubernetes command line tool.

## About this task

The log data collected and stored by Logstash for your IBM Cloud Private environment can be configured to be forwarded to IBM Cloud App Management as event information and then correlated into incidents.

## Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **Logstash** tile and click **Configure**.
4. Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.
5. Click **Save**.
6. Modify the default Logstash configuration in IBM Cloud Private to add IBM Cloud App Management as a receiver. To do this, edit the Logstash pipeline ConfigMap to add the webhook URL in the output section as follows:

- a) Load the ConfigMap into a file using the following command:

```
kubectl get configmaps logstash-pipeline --namespace=kube-system -o yaml > logstash-pipeline.yaml
```

**Note:** The default Logstash deployment ConfigMap name in IBM Cloud Private is logstash-pipeline in the kube-system namespace. If your IBM Cloud Private logging uses a different Logstash deployment, modify the ConfigMap name and namespace as required for that deployment.

- b) Edit the logstash-pipeline.yaml file and add an HTTP section to specify IBM Cloud App Management as a destination using the generated webhook URL. Paste the webhook URL into the **url** field:

```
output {
  elasticsearch {
    index => "logstash-%{+YYYY.MM.dd}"
    hosts => "elasticsearch:9200"
  }
  http {
    url => "<Cloud_Event_Management_webhook_URL>"
    format => "json"
    http_method => "post"
    pool_max_per_route => "5"
  }
}
```

**Note:** The pool\_max\_per\_route value is set to 5 by default. It limits the number of concurrent connections to IBM Cloud App Management to avoid data overload from Logstash. You can modify this setting as required.

- c) Save the file, and replace the ConfigMap using the following command:  

```
kubectl --namespace kube-system replace -f logstash-pipeline.yaml
```
- d) Check the update is complete at [https://<icp\\_master\\_ip\\_address>:8443/console/configuration/configmaps/kube-system/logstash-pipeline](https://<icp_master_ip_address>:8443/console/configuration/configmaps/kube-system/logstash-pipeline)

**Note:** It can take up to a minute for the configuration changes to take affect.

7. To start receiving log data from Logstash, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.


## Configuring Microsoft Azure as an event source

Microsoft Azure provides monitoring services for Azure resources. You can set up an integration with Cloud Event Management to receive alert information from Microsoft Azure. The Microsoft Azure integration is only available in IBM Cloud App Management, Advanced.

### About this task

Using a webhook URL, alerts generated by Microsoft Azure monitoring are sent to the IBM Cloud App Management service as events.

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **Microsoft Azure** tile and click **Configure**.
4. Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.
5. Click **Save**.
6. Log in to your Microsoft Azure account at <https://portal.azure.com/>.
7. Go to the **Dashboard** and select the resource you want event information from. Click the resource name.
8. Go to **MONITORING** in the navigation menu and click **Alert rules**.
9. Click **Add alert** at the top of the page.
10. Set up the rule as follows:
  - a) Enter a name for the rule and add a description.
  - b) Select the metric that you want this alert rule to monitor for the selected resource.
  - c) Set a condition and enter a threshold value for the metric. When the threshold value for the set condition is reached, an alert is generated and sent as an event to IBM Cloud App Management.
  - d) Select the time period to monitor the metric data.
  - e) Optional: Set up email notification.
  - f) Paste the webhook URL into the **Webhook** field. This is the generated URL provided by IBM Cloud App Management.
  - g) Click **OK**.
11. To start receiving alert information from Microsoft Azure, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.

## Configuring Nagios XI as an event source

Nagios XI provides network monitoring products. You can set up an in integration with IBM Cloud App Management to receive alert information from Nagios XI products. The Nagios XI integration is only available in IBM Cloud App Management, Advanced.

### About this task

Using a package of configuration files provided by IBM Cloud App Management, you set up an integration with Nagios XI. The alerts generated by Nagios XI are sent to the IBM Cloud App Management service as events.

**Note:** IBM Cloud App Management supports integration with the server monitoring and web monitoring components of the Nagios XI product.

## Procedure

1. Ensure that the Nagios Plugins are installed into your instance of Nagios XI. Depending on how the plugins are controlled, you can check their status as follows:
  - If you use **xinetd** for controlling the plugins: **service xinetd status**
  - If you use a dedicated daemon for controlling the plugins: **service nrpe status**
2. Click **Integrations** on the IBM Cloud App Management **Administration** page.
3. Click **Configure an integration**.
4. Go to the **Nagios XI** tile and click **Configure**.
5. Enter a name for the integration.
6. Click **Download file** to download the `nagios-cem.zip` file. The compressed file contains three files to set up the integration with IBM Cloud App Management:
  - The file `cem.cfg` needs to be imported into Nagios XI.
  - The file `nagios-cem-webhook.sh` includes the unique webhook URL generated for this integration.
  - The file `import-cem.sh` copies the `cem.cfg` and `nagios-cem-webhook.sh` files to Nagios XI destination directory.
7. Click **Save** to save the integration in IBM Cloud App Management.
8. Extract the files to any directory, and copy the files to the Nagios XI server.
9. Run the `import-cem.sh` command to copy the `cem.cfg` and `nagios-cem-webhook.sh` files to the correct Nagios XI destination directory.

For example, if you are logged in as a non-root user, run the command as follows to ensure it runs as root and copies the files as required: `sudo bash ./import-cem.sh`.

10. Log in to the Nagios XI UI as an administrator, and use the **Core Config Manager** to import the `cem.cfg` file:
  - a) Go to **Configure** in the menu bar at the top of the window and select **Core Config Manager** from the list.
  - b) Select **Tools > Import Config Files** from the menu on the left side of the window.
  - c) Select `cem.cfg` and click **Import**.
11. Enable the environment variable macro:
  - a) In **Core Config Manager**, select **CCM Admin > Core Configs** from the menu on the left side of the window.
  - b) On the **General** tab enter 1 for the `enable_environment_macros` parameter.
  - c) Click **Save Changes**.
12. Ensure the `cemwebhook` contact is added to the set of hosts and services you monitor:

**Note:** Remember to enable the `cemwebhook` contact when setting up a source to monitor. To enable the `cemwebhook` contact for the host and all services for that host, ensure you select **CEM Webhook-Contact** under **Send Alert notification To** in Step 4 of the Configuration Wizard.

To check that `cemwebhook` is among the contacts included in alerts for a host:

- a) In **Core Config Manager**, select **Monitoring > Hosts** from the menu on the left side of the window.
- b) Click a host name to edit its settings.
- c) Click the **Alert Settings** tab and then click **Manage Contacts**.
- d) Ensure that `cemwebhook` is in the **Assigned** column. If not, then select it and click **Add Selected**.
- e) Click **Close** and then **Save**.

**Note:** This example is for checking host settings, but the same steps can be followed to check services.

13. Change the command type for the **notify-cem-host** and **notify-cem-service** commands:

- a) In **Core Config Manager**, select **Commands** > **\_Commands** from the menu on the left side of the window.
  - b) Locate and click **notify-cem-host** to edit its settings.
  - c) Select **misc command** from the **Command Type** list.
  - d) Click **Save**.
  - e) Repeat for **notify-cem-service**.
14. Select **Quick Tools** > **Apply Configuration** from the menu on the left side of the window and click **Apply Configuration**.
  15. To start receiving alert information from Nagios XI, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.


## Configuring Netcool/OMNIBus as an event source

Tivoli Netcool/OMNIBus is a service level management (SLM) system that delivers real-time, centralized monitoring of complex networks and IT domains. You can integrate with an existing on-premises installation of Netcool/OMNIBus to receive event information about the resources it monitors.

### About this task

Events from Netcool/OMNIBus are received through a gateway you install and configure.

### Procedure

1. Click **Users and Groups** on the Cloud App Management **Administration** user interface.
2. Click **Configure an integration**.
3. Go to the **Netcool/OMNIBus** tile and click **Configure**.
4. Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.
5. Click **Save**.
6. Install and configure the Netcool/OMNIBus Gateway for Cloud App Management:
  - a) Download the gateway by clicking the **Download the gateway** link on the UI.
  - b) Copy the package to the server you want to install it on.
  - c) Extract the package and follow the instructions in the README file to install and configure the Gateway.
  - d) As part of configuring, ensure you edit the `G_CEM.props` file and add the generated webhook URL to the `Gate.CEM.WebhookURL` property.
  - e) Save the `G_CEM.props` file.
7. To start receiving events from Netcool/OMNIBus, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.

### What to do next

For more information about Netcool/OMNIBus, see <https://www.ibm.com/support/knowledgecenter/en/SSSTQ/landingpage/NetcoolOMNIBus.html>.

## Configuring New Relic as an event source


New Relic monitors mobile and web applications in real-time, helping users diagnose and fix application performance problems. You can receive New Relic alerts through the incoming webhooks of the IBM

Cloud App Management service. The New Relic integration is only available in IBM Cloud App Management, Advanced.

### About this task

You can configure integration with both New Relic Legacy or New Relic Alerts systems. Both configuration procedures are documented here. The first step is to generate the webhook URL within IBM Cloud App Management.

### Procedure

1. Generate an incoming webhook for New Relic:
  - a) Click **Integrations** on the IBM Cloud App Management **Administration** page.
  - b) Click **Configure an integration**.
  - c) Depending on the version you use, go to the **New Relic Legacy** or **New Relic Alerts** tile, and click **Configure**.
  - d) Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.
  - e) Click **Save**.
2. Use the incoming webhook to:
  - [“Configure New Relic Legacy” on page 706](#) as source.
  - [“Configure New Relic Alerts” on page 707](#) as source.

### Configure New Relic Legacy

Configure integration with New Relic Legacy.

### About this task

Configure New Relic Legacy as source:

### Procedure

1. Generate an incoming webhook as described in [“1” on page 706](#).
2. Log in to New Relic at <https://rpm.newrelic.com/> as an administrator.
3. From the New Relic menu bar, select **Alerts > Channels and groups**.
4. In the **Channel details** section, click **Create channel > Webhook**.
5. Enter a name for the channel and paste the incoming webhook URL into the **Webhook URL** field. This is the generated URL provided by IBM Cloud App Management. Add an optional description.
6. Select your **Notification level**.
7. Click **Integrate with Webhooks**.
8. Associate the webhook channel with all of the New Relic policies that you want to receive events from. For more information about associating channels with policies, see the New Relic documentation at <https://docs.newrelic.com/docs/alerts/new-relic-alerts/managing-notification-channels/add-or-remove-policy-channels>.
9. To start receiving events from New Relic, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.

## Configure New Relic Alerts

Configure integration with New Relic Alerts.

### About this task

Configure New Relic Alerts as source:

### Procedure

1. Generate an incoming webhook as described in “1” on page 706.
2. Log in to New Relic at <https://alerts.newrelic.com/> as an administrator.
3. From the New Relic menu bar, select **Alerts > Notification channels**.
4. Click **New notification channel**.
5. In the **Channel details** section, select Webhook for channel type.
6. Enter a name for the channel and paste the webhook URL into the **Base URL** field. This is the generated URL provided by Cloud Event Management.
7. Click **Create channel**.
8. Associate the webhook channel with all of the New Relic policies that you want to receive events from. For more information about associating channels with policies, see the New Relic documentation at <https://docs.newrelic.com/docs/alerts/new-relic-alerts/managing-notification-channels/add-or-remove-policy-channels>.
9. Ensure you set the incident preference to **By condition and entity**. This is required to send notifications to IBM Cloud App Management every time a policy violation occurs. IBM Cloud App Management uses this information to accurately correlate events into incidents, and clear them when applicable.
  - a) From the New Relic menu bar, select **Alerts > Alert policies**.
  - b) Select your alert policy and click **Incident preference**.
  - c) Select **By condition and entity**, and click **Save**.
  - d) Repeat for each alert policy that sends notifications to IBM Cloud App Management.

For more information about incident preferences in New Relic, see <https://docs.newrelic.com/docs/alerts/new-relic-alerts/configuring-alert-policies/specify-when-new-relic-creates-incidents>.
10. To start receiving events from New Relic, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.

## Configuring Pingdom as an event source


Pingdom provides web performance and availability monitoring. You can set up an integration with IBM Cloud App Management to receive alert information from Pingdom. The Pingdom integration is only available in IBM Cloud App Management, Advanced.

### About this task

Using a webhook URL, you set up an integration with Pingdom, and associate the integration with the uptime and transaction checks. The alerts generated by the checks are sent to the IBM Cloud App Management service as events.

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **Pingdom** tile and click **Configure**.

4. Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.
5. Click **Save**.
6. Log in to your account at <https://my.pingdom.com/>.
7. Set up the integration:
  - a) Select **Integrations > Integrations**.
  - b) Click **Add new** in the upper-right corner of the window.
  - c) Ensure **Webhook** is selected from the **Type** list.
  - d) In the **Name** field, enter a name for the integration.
  - e) In the **URL** field, paste the webhook URL from IBM Cloud App Management.
  - f) Ensure the **Active** check box is selected.
  - g) Click **Save integration**.

**Tip:** For more information about setting up webhook integrations in Pingdom, see <https://help.pingdom.com/hc/en-us/articles/207081599>.
8. Enable the integration for the checks you want to receive alert information from:
  - a) Go to <https://my.pingdom.com/dashboard>.
  - b) Select **Monitoring > Uptime**.
  - c) Open a check, and select the check box next to your webhook integration. This enables the posting of alerts to the URL when, for example, a site goes down.

**Tip:** If you don't have checks set up, you can add them by clicking **Add new** in the upper-right corner of the window. For more information about checks in Pingdom and how to set them up, see <https://help.pingdom.com/hc/en-us/articles/203749792-What-is-a-check->.
  - d) Repeat the steps for each check you want to receive alert information from.
9. To start receiving alert information from the Pingdom checks, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.


## Configuring Prometheus as an event source

Prometheus is an open-source systems monitoring and alerting toolkit. You can set up an integration with IBM Cloud App Management to receive alert information from Prometheus. The Prometheus integration is only available in IBM Cloud App Management, Advanced.

### About this task

Using an incoming webhook URL, configure your Prometheus instance to route alerts to IBM Cloud App Management, and define alerting rules in your Prometheus Alertmanager configuration file.

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **Prometheus** tile and click **Configure**.
4. Enter a name for the integration and click  **Copy** to add the generated webhook URL to the clipboard. Ensure you save the generated webhook to make it available later in the configuration process. For example, you can save it to a file.
5. Click **Save**.
6. Set up the integration in Prometheus as follows:



**Note:** If you want to receive event information from Prometheus included in IBM Cloud Private, the following steps are different. See step “7” on page 710 for details about how to configure Prometheus included in IBM Cloud Private.

- a) Ensure you have the Prometheus Alertmanager installed as described in <https://github.com/prometheus/alertmanager#installation>.
- b) Configure the Alertmanager to send alert information from Prometheus to Cloud Event Management. Edit the `alertmanagerFiles` section of your Alertmanager configuration file to add the generated webhook from IBM Cloud App Management as a receiver. Paste the webhook into the `url` field. In addition, set the `send_resolved` value to `true`.  
For example:

```
alertmanagerFiles:
  alertmanager.yml: |-
    global:
      resolve_timeout: 20s

    receivers:
      - name: 'webhook'
        webhook_configs:
          - send_resolved: true
            url: 'https://myeventsourcesource.mybluemix.net/webhook/prometheus/omaasdev/63234831-4389-480f-8035-bc293b4e05fe/1pA0lWhP09t9_FhPLxNyKrGuglYBnHPa1MXbx4otg3Y'

    route:
      group_wait: 10s
      group_interval: 5m
      receiver: webhook
      repeat_interval: 3h
```

For more information about Alertmanager configuration files, see <https://prometheus.io/docs/alerting/configuration/>.

- c) Edit the `serverFiles` section of your Alertmanager configuration file to define your alerting rules. You must provide at least the following fields for each alert: severity, summary, description, and type. Severity must be one of the following values:
- indeterminate
  - information
  - warning
  - minor
  - major
  - critical
  - fatal

The alerting rules syntax is different depending on the version of Prometheus you are using.

If you are using Prometheus version 1.8, see the following example for alerting rules:

```
serverFiles:
  rules: ""
  alerts: |-
    # Rules for Node
    ALERT high_node_load
    IF node_load1 > 20
    FOR 10s
    LABELS { severity = "critical" }
    ANNOTATIONS {
      # summary defines the status if the condition is met
      summary = "Node usage exceeded threshold",
      # description reports the situation of event
      description = "Instance {{ $labels.instance }} Job {{ $labels.job }},
        Node load {{ $value }}",
      # type defines the type of the resource causing the event
      type = "Server",
    }
  }
```

```

ALERT high_memory_usage
IF (( node_memory_MemTotal - node_memory_MemFree ) / node_memory_MemTotal) *
100 > 100
FOR 10s
LABELS { severity = "warning" }
ANNOTATIONS {
# summary defines the status if the condition is met
summary = "Memory usage exceeded threshold",
# description reports the situation of event
description = "Instance {{ $labels.instance }} , Job {{ $labels.job }} ,
Memory usage {{ humanize $value }}%",
# type defines the type of the resource causing the event
type = "Server",
}

ALERT high_storage_usage
IF (node_filesystem_size{fstype="ext4"} -
node_filesystem_free{fstype="ext4"}) /
node_filesystem_size{fstype="ext4"} * 100 > 90
FOR 10s
LABELS { severity = "warning" }
ANNOTATIONS {
# summary defines the status if the condition is met
summary = "Storage usage exceeded threshold",
# description reports the situation of event
description = "Instance {{ $labels.instance }} , Job {{ $labels.job }} ,
Storage usage {{ humanize $value }}%",
# type defines the type of the resource causing the event
type = "Storage",
}

```

If you are using Prometheus version 2.0 or later, see the following example for alerting rules:

```

- alert: high_cpu_load
  expr: node_load1 > 60
  for: 30s
  labels:
    severity: critical
  annotations:
    description: Docker host is under high load, the avg load 1m is at {{ $value }}.
      Reported by instance {{ $labels.instance }} of job {{ $labels.job }}.
    summary: Server under high load
    type: Server
- alert: high_memory_load
  expr: (sum(node_memory_MemTotal) - sum(node_memory_MemFree + node_memory_Buffers
    + node_memory_Cached)) / sum(node_memory_MemTotal) * 100 > 85
  for: 30s
  labels:
    severity: warning
  annotations:
    description: Docker host memory usage is {{ humanize $value }}%. Reported by
      instance {{ $labels.instance }} of job {{ $labels.job }}.
    summary: Server memory is almost full
    type: Server
- alert: high_storage_load
  expr: (node_filesystem_size{fstype="aufs"} - node_filesystem_free{fstype="aufs"})
    / node_filesystem_size{fstype="aufs"} * 100 > 85
  for: 30s
  labels:
    severity: warning
  annotations:
    description: Docker host storage usage is {{ humanize $value }}%. Reported by
      instance {{ $labels.instance }} of job {{ $labels.job }}.
    summary: Server storage is almost full
    type: Server


```

**Tip:** For more information about Prometheus alerting rules, see [https://prometheus.io/docs/prometheus/2.2/configuration/alerting\\_rules/](https://prometheus.io/docs/prometheus/2.2/configuration/alerting_rules/).


d) Save and close the file.

7. If you want to receive event information from Prometheus included in IBM Cloud Private, set up the integration using the IBM Cloud Private UI as follows:

- a) Log in to your IBM Cloud Private host. From the navigation menu, click **Configuration** > **ConfigMaps**.
- b) Search for alert to list the ConfigMaps for the Prometheus Alertmanager and alerting rules.
- c) Configure the Alertmanager to send alert information from Prometheus in IBM Cloud Private to IBM Cloud App Management. Edit the monitoring-prometheus-alertmanager ConfigMap by

clicking  and Edit. Add the generated webhook from IBM Cloud App Management as a receiver. Paste the webhook into the url: field. In addition, set the send\_resolved value to true. You can also click **Create resource**, add the following Alertmanager configuration, paste the webhook from IBM Cloud App Management into the url: field, and click **Create**. This will overwrite your settings in monitoring-prometheus-alertmanager (note that this example also includes a Slack channel configuration):

```
apiVersion: v1
kind: ConfigMap
metadata:
  labels:
    app: monitoring-prometheus
    component: alertmanager
    name: monitoring-prometheus-alertmanager
    namespace: kube-system
data:
  alertmanager.yml: |-
    global:
      resolve_timeout: 20s
      slack_api_url: 'https://hooks.slack.com/services/xxx/yyy/zzz'
    route:
      receiver: webhook
      group_by: [alertname, instance, severity]
      group_wait: 10s
      group_interval: 10s
      repeat_interval: 1m
      routes:
        - receiver: webhook
          continue: true
        - receiver: slack_alerts
          continue: true
      receivers:
        - name: webhook
          webhook_configs:
            - send_resolved: true
              url: 'https://<webhook_url_from_Cloud_Event_Managent.net/webhook/
prometheus/xxx/yyy/zzz'
        - name: slack_alerts
          slack_configs:
            - send_resolved: false
              channel: '#ibmcloudprivate'
              text: 'Nodes: {{ range .Alerts }}{{ .Labels.instance }} {{ end }}
---- Summary: {{ .CommonAnnotations.summary }}
---- Description: {{ .CommonAnnotations.description }}
---- https://9.30.189.183:8443/prometheus/alerts '
```

- d) Edit the monitoring-prometheus-alertrules ConfigMap to define your alerting rules. Click  and **Edit**.

You must provide at least the following fields for each alert: severity, summary, description, and type. Severity must be one of the following values:

- indeterminate
- information
- warning
- minor
- major
- critical
- fatal

You can also click **Create resource**, add the following alerting rules, and click **Create**. This will overwrite your settings in monitoring-prometheus-alertrules:

```
apiVersion: v1
kind: ConfigMap
metadata:
  labels:
    app: monitoring-prometheus
    component: prometheus
  name: monitoring-prometheus-alertrules
  namespace: kube-system
data:
  sample.rules: |-
    groups:
    - name: alert.rules
      rules:
      - alert: high_cpu_load
        expr: node_load1 > 5
        for: 10s
        labels:
          severity: critical
        annotations:
          description: Docker host is under high load, the avg load 1m is at {{ $value }}.
            Reported by instance {{ $labels.instance }} of job {{ $labels.job }}.
          summary: Server under high load
      - alert: high_memory_load
        expr: (sum(node_memory_MemTotal) - sum(node_memory_MemFree + node_memory_Buffers
          + node_memory_Cached)) / sum(node_memory_MemTotal) * 100 > 85
        for: 30s
        labels:
          severity: warning
        annotations:
          description: Docker host memory usage is {{ humanize $value }}%. Reported by
            instance {{ $labels.instance }} of job {{ $labels.job }}.
          summary: Server memory is almost full
      - alert: high_storage_load
        expr: (node_filesystem_size{fstype="aufs"} - node_filesystem_free{fstype="aufs"})
          / node_filesystem_size{fstype="aufs"} * 100 > 15
        for: 30s
        labels:
          severity: warning
        annotations:
          description: Docker host storage usage is {{ humanize $value }}%. Reported by
            instance {{ $labels.instance }} of job {{ $labels.job }}.
          summary: Server storage is almost full
```

- e) Optional: To check that you have set up Prometheus in IBM Cloud Private to send event information: from the navigation menu, click **Platform > Alerting**, and click the **Status** tab; check that your settings are available in the **Config** section.
8. To start receiving alert information from Prometheus, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.

## Configuring SolarWinds Orion as an event source

The SolarWinds Orion platform provides network and system management products. You can set up an integration with IBM Cloud App Management to receive alert information from SolarWinds Orion. The SolarWinds integration is only available in IBM Cloud App Management, Advanced.

### About this task

Using an XML file, you set up an integration with SolarWinds Orion, and define trigger and reset actions for alerts. The alerts generated by SolarWinds Orion are sent to the IBM Cloud App Management service as events.

**Note:** IBM Cloud App Management supports integration with the Network Performance Monitor and Server and Application Monitor products of the SolarWinds Orion platform.

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.

3. Go to the **SolarWinds** tile and click **Configure**.
4. Enter a name for the integration.
5. Click **Download file** to download the `send-alert-cem.xml` file. This file contains the settings required for the integration with IBM Cloud App Management, including the webhook URL.

**Note:** If you edit the integration later and click to download the file again, the current integration will no longer be valid. You will need to set up the integration again.

6. Click **Save** to save the integration in IBM Cloud App Management.
7. Upload the XML file to the **Alert Manager** in SolarWinds Orion:
  - a) Log in to your SolarWinds Orion account as an administrator.
  - b) Go to **ALERTS & ACTIVITY** in the menu bar at the top of the window and select **Alerts** from the list.
  - c) Click **Manage alerts**.
  - d) Go to **EXPORT/IMPORT** in the menu bar at the top of the window and select **Import Alert** from the list.
  - e) Upload the `send-alert-cem.xml` you downloaded earlier from IBM Cloud App Management.

**Note:** A new alert called **Notify CEM - timestamp** is created, together with the associated trigger and reset actions **Post Problem Event to CEM - timestamp** and **Post Resolution Events to CEM - timestamp**, where *timestamp* is in the UTC format. The **Notify CEM** alert contains settings for the integration between IBM Cloud App Management and SolarWinds. It is disabled by default and is not intended to be enabled.

8. Define trigger and reset actions for the alerts you want IBM Cloud App Management to receive event information from:
  - a) In **Alert Manager**, click the alert you want to edit, and go to the **TRIGGER ACTIONS** tab.
  - b) Click the **Assign Action(s)** button.
  - c) Select the **Post Problem Event to CEM - timestamp** check box and click **ASSIGN**.
  - d) Click **Next** to go to the **RESET ACTION** tab.
  - e) Click the **Assign Action(s)** button.
  - f) Select the **Post Resolution Events to CEM - timestamp** check box and click **ASSIGN**.
  - g) Click **Next** and then click **Submit**.



**Attention:** If you create more than one SolarWinds integration instance, ensure you select the right trigger and reset actions for each integration. For example, for your first integration select **Post Problem Event to CEM - timestamp1** and **Post Resolution Events to CEM - timestamp1**, while for your second integration select **Post Problem Event to CEM - timestamp2** and **Post Resolution Events to CEM - timestamp2**.

If you have more than one SolarWinds integration instance, you can find out which integration sent specific events by checking the detailed event information. Go to the incident, click the **Events** link in the incident card, expand the event, and click the **See more info** button. See the **Event source** table for details about the system that sent the event, such as the event source name and type.

**Tip:** You can also define the trigger and reset actions for more than one alert at the same time. For the trigger action, select the check box for the alerts and select **Assign Trigger Action** from the **ASSIGN ACTION** list. Then select the **Post Problem Event to CEM - timestamp** check box and click **ASSIGN**. For the reset action, select the check box for the same alerts and select **Assign Reset Action** from the **ASSIGN ACTION** list. Then select the **Post Resolution Events to CEM - timestamp** check box and click **ASSIGN**.

**Note:** IBM Cloud App Management supports Out-Of-The-Box Alerts (OOTBA) for the following common objects in SolarWinds:

- Application

- Component
- Group
- Interface
- Node
- Volume

You can check the object type of each alert in **Alert Manager** by looking at the **Property to Monitor** column for an alert.

If you enable an unsupported alert type, event information might still be sent to IBM Cloud App Management, but the event title will state "Unsupported SolarWinds object".

9. To enable the alert, set **Enabled (On/Off)** to **On** in the appropriate rows for the alerts you want to receive event information from.
10. To start receiving alert information from the SolarWinds Orion triggers and reset actions, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.

## Configuring Splunk Enterprise as an event source

Splunk Enterprise is an on-premises version of Splunk that you can use to monitor and analyze machine data from various sources. You can set up an integration with IBM Cloud App Management to receive alert information from Splunk Enterprise. The Splunk Enterprise integration is only available in IBM Cloud App Management, Advanced.

### About this task

Using a package of installation and configuration files provided by IBM Cloud App Management, you set up an integration with Splunk Enterprise. The alerts generated by Splunk Enterprise are sent to the IBM Cloud App Management service as events.

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **Splunk Enterprise** tile and click **Configure**.
4. Enter a name for the integration.
5. Click **Download file** to download and decompress the `ibm-cem-splunk.zip` file. The compressed file contains the `savedsearches.conf` file for both the Unix and Windows systems, and the `ibm-cem-alert.zip` file which contains the file for installing the Splunk App for IBM Cloud App Management.
  - `splunk_app_for_nix/local/savedsearches.conf`
  - `splunk_app_windows_infrastructure/local/savedsearches.conf`
  - `ibm-cem-alert.zip`
6. Install the Splunk App using the `ibm-cem-alert.zip` file.
  - a) Log in to your Splunk Enterprise browser UI as an administrator.
  - b) Select **App** then click **Manage Apps**.
  - c) Click **Install app from file**.
  - d) Click **Browse** to locate the `ibm-cem-alert.zip` file.
  - e) Click **Upload**.
7. Log in to your Splunk Enterprise server host and copy the `savedsearches.conf` file to `$SPLUNK_HOME/etc/apps/<app_name>/local`.

Linux:

```
sudo cp ibm-cem-splunk/splunk_app_for_nix/local/savedsearches.conf
  $SPLUNK_HOME/etc/apps/splunk_app_for_nix/local/savedsearches.conf
```

Windows:

```
copy ibm-cem-splunk\splunk_app_windows_infrastructure\local\savedsearches.conf
  %SPLUNK_HOME%\etc\apps\splunk_app_windows_infrastructure\local
```

**Important:** If you already have an existing Splunk app installed, then you already have settings defined in a `savedsearches.conf` file. Merge your existing `savedsearches.conf` file with the one downloaded from IBM Cloud App Management. You can merge the files manually, or use the Splunk Enterprise browser UI by clicking the **Alerts** tab at the top, expanding the selected alert section, clicking **Edit** > **Edit Alerts**, and editing the fields under section **IBM Cloud Event Management Alert**. You can use the `savedsearches.conf` file to check the mapping for the values of the fields.

- Restart the Splunk Enterprise instance to ensure the new alerts are available.

Unix:

```
sudo $SPLUNK_HOME/bin/splunk restart
```

Windows:

```
%SPLUNK_HOME%\bin\splunk.exe restart
```

- Log in to the Splunk Enterprise UI as an administrator and check that the alerts defined in `savedsearches.conf` are available:

For Unix systems, go to **Search & Reporting** > **Splunk App for Unix** > **Core Views** > **Alerts**.

For Windows systems, go to **Search & Reporting** > **Splunk App for Windows Infrastructure** > **Core Views** > **Alerts**.

**Note:** If you modify the trigger conditions for the alerts, ensure you do not set a trigger interval that is too frequent. For example, if you set the **Edit** > **Edit Alerts** > **Trigger Conditions** to trigger an alert once every minute when the result count is greater than 0, the resulting number of events can overload IBM Cloud App Management. To limit the trigger frequency, set the **greater than** value to a higher number than 0, and set it to be triggered 5 times in every hour, for example. You can also use the **Throttle** option to suspend the triggering of events for a set period after an event is triggered.

- Optional: To receive resolution events from Splunk Enterprise, add the `resolution:true` value to the `action.ibm_cem_alert.param.cem_custom` parameter in the `savedsearches.conf` file, for example:

```
# Example
### Automation mapping for IO Utilization Exceeds Threshold Alert
## using IBM Event Management custom webhook alert
[IO_Utilization_Exceeds_Threshold]
action.ibm_cem_alert = 1
action.ibm_cem_alert.param.cem_custom = statusOrThreshold:$result.bandwidth_util
$,resolution:true
action.ibm_cem_alert.param.cem_event_type = $name$
action.ibm_cem_alert.param.cem_resource_name = $result.host$
action.ibm_cem_alert.param.cem_resource_type = Server
action.ibm_cem_alert.param.cem_severity = Major
action.ibm_cem_alert.param.cem_summary = $result.host$: IO utilization exceeds
  $bandwidth_util$ threshold
action.ibm_cem_alert.param.cem_webhook = {{WEBHOOK_URL}}/{{WEBHOOK_USER}}/
  {{WEBHOOK_PASSWORD}}
disabled = 0
```

**Tip:** You can also add the resolution setting using the UI. Open **Edit** > **Edit Alerts** under section **IBM Cloud Event Management Alert**, and add `resolution:true` to the **Additional mapping (optional)** field.

- Click **Save** to save the integration in IBM Cloud App Management.

12. To start receiving alert notifications from Splunk Enterprise, ensure that **Enable event management from this source** is set to **On** in IBM Cloud App Management.

## Configuring IBM UrbanCode Deploy as an event source

You can set up an integration with IBM Cloud App Management to receive notifications created by IBM UrbanCode Deploy. IBM UrbanCode Deploy is a tool for automating application deployments through your environments. It facilitates rapid feedback and continuous delivery in agile development, while providing the audit trails, versioning, and approvals needed in production. Emails are sent to IBM Cloud App Management as events. The Urban Code Deploy integration is only available in IBM Cloud App Management, Advanced.

### Before you begin

You must have a docker account.

### About this task

IBM UrbanCode Deploy sends email notifications when user-defined trigger events occur on the server. You must configure the email probe container to retrieve emails from the email account and perform the normalization. After the normalization, the probe will send the events to IBM Cloud App Management.

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click **Configure an integration**.
3. Go to the **IBM UrbanCode Deploy** tile and click **Configure**.
4. Enter a name for the integration.
5. Click **Download file** to download and decompress the `email-probe-package.zip` file.
6. Extract the package into a docker environment where docker and docker compose are installed.
7. Grant execution rights to `integration.sh`, for example `chmod 755 integration.sh`.
8. Go to <https://store.docker.com/images/ibm-netcool-probe-email> to read the description and then click **Proceed to checkout** on the right of the page. Enter the required contact information and click **Get Content**.
9. Run `docker login` in your docker environment.
10. Uncomment `LICENSE=accept` in `probe.env` to accept the license agreement.
11. Update `probe.env` to populate `EMAIL_SERVER_HOSTNAME`, `USERNAME`, `PASSWORD`, and `FILTER`.
  - `EMAIL_SERVER_HOSTNAME` is used to specify the email server host name, such as `gmail.com`.
  - `USERNAME` is used to specify the user name of the email account.
  - `PASSWORD` is used to specify the plain password to access the email account. The plain password will be encrypted when the container is running. **Note:** do not set a password that starts with `ENCRYPTED` as this keyword is used to determine whether it is a plain or encrypted password.
  - `FILTER` is used to specify the UCD sender email address. The sender email address can be found in **UCD Home > Settings > System Settings > Mail Server Settings > Mail Server Sender Address**.
  - Optionally, you can update `POLLINTERVAL` to specify the frequency in seconds for the probe to retrieve new emails. The default value is 600 seconds.
12. Run `docker-compose up` to start the probe.

### Note:

- a. The probe only connects to the IMAP mail server over a TLS connection so the email server must have a valid certificate.
- b. The probe only supports the default UCD notification template.
- c. The probe deletes emails from the mail server after retrieving them.
- d. For the probe to run smoothly, avoid updating other probe properties in `email.props`.



## Unsupported Events

There are four mandatory IBM Cloud App Management fields required to publish an event in IBM Cloud App Management. The attributes are **Resource Name**, **Summary**, **Event Type** and **Severity**. If an Unknown - *<cause>* error is displayed for any of these fields, you will need to update the [UCD email notification template](#). This might happen if you have used a custom notification template.

The following table contains some error messages, possible causes and resolutions.

CEM Field	Messages	Possible Causes	Resolutions
Summary	Unknown - Missing the Subject field in UCD email.	Missing the Subject field in UCD email.	Update the notification template to add Subject field.
Resource Name	Unknown - Missing the expected format of Application name in UCD email.	The Application field in the UCD email is not following the format in default email notification template.	Follow the exact format used in the default notification templates.
Resource Name	Unknown - Missing the expected format of Process name in UCD email.	The Process field in the UCD email is not following the format in default email notification template.	Follow the exact format used in the default notification templates.
Resource Name	Unknown - Missing the Application or Process name in UCD email.	Missing the Application or Process field in UCD email.	Update the notification template to add Application or Process field.
Event Type	Unknown - Missing the keyword of Process or Approval at the Subject field in UCD email to indicate the event type.	Missing the keyword of Process or Approval at the Subject field in UCD email to indicate the event type.	Update the notification template to add keyword of Process or Approval at the Subject field.

If you feel that the current webhook URL for the email probe has been compromised in some way, you can download the email probe zip file again to regenerate the webhook. This invalidates the existing webhook URL and replaces it with a new one. In this scenario, you must repeat the configuration steps to save the zip file in a docker environment and rerun docker compose to start the email probe with new webhook.

## Configuring outgoing event destinations

Integrate with the tools and systems to which you want to send events and data from Cloud App Management. For example, you can integrate with IBM Tivoli® Netcool/OMNIbus to forward events to the event manager, or integrate with Stride to forward events and metrics to your custom applications.

## Sending incident details to Alert Notification


When using an IBM Cloud Private environment, you can send incident details to your Alert Notification service. The Alert Notification service can then notify the right teams about the incidents as required. You can set up this outgoing integration using the Alert Notification API. The Alert Notification integration is only available in IBM Cloud App Management, Advanced.

### Before you begin

Ensure you meet the following prerequisites:

- Ensure you have an Alert Notification service set up. For more information, see [IBM Alert Notification Documentation](#).

### Procedure

1. Go to your Alert Notification service and generate an API key as described in the [Managing API keys](#) topic in the IBM Alert Notification knowledge center. Make a note of the API key name and the API key password.
2. Click **Integrations** on the **Administration** page.
3. Click **Configure an integration**, and click the **Outgoing** tab.
4. Go to the **Alert Notification** tile and click **Configure**.
5. Enter a name for the integration.
6. Enter the URL for your Alert Notification API. You can check your API URL by clicking **Manage API Keys** on the navigation menu in Alert Notification, and clicking the  icon.
7. Enter the API key name and API key password in the **Enter your Alert notification API credentials** section. These are the values you generated in Alert Notification.
8. Ensure that **Enable integration** is set to **On** in Cloud Event Management.
9. Click **Save**.
10. Set up an incident policy where you set your Alert Notification integration as a recipient of notifications. For more information, see [“Managing incident policies” on page 742](#).

[Creating custom notification templates](#)

[Context data](#)

## Configuring Netcool/OMNIBus as an outgoing incident source

You can set up Cloud App Management to send incident details to your on-premises Netcool/OMNIBus installation. You can set up this outgoing integration by using the Netcool/OMNIBus for Message Bus.

### Before you begin

Ensure that you meet the following prerequisites:

- Ensure that you have Cloud App Management set up in your IBM Cloud Private environment.
- Ensure that you have an existing Netcool/OMNIBus deployment. For more information, see [Installing and updating Netcool/OMNIBus](#) in the IBM Knowledge Center.
- Ensure you have the Netcool/OMNIBus for Message Bus installed. To download the probe, see [IBM Tivoli Netcool/OMNIBus probes and gateways](#) in the IBM Knowledge Center. To install the probe, see [Installing probes and gateways on Netcool/OMNIBus V8.1](#) in the IBM Knowledge Center.

### Procedure

1. Click **Integrations** on the Cloud App Management console **Administration** page.
2. Click the **Outgoing** tab, and click **Configure an integration**.
3. Go to the **Netcool/OMNIBus** tile and click **Configure**.
4. Enter a name for the integration.
5. Click **Download file** to download the `ibm-cem-noi-config.zip` file, and save it. The compressed file contains configuration settings for the Message Bus Probe and the `apm.sql` file used to update the Netcool/OMNIBus ObjectServer database.
6. Copy and extract the `ibm-cem-noi-config.zip` file to the Netcool/OMNIBus ObjectServer system.
7. Update the Netcool/OMNIBus ObjectServer database schema by loading the `apm.sql` file into the database:

```
$OMNIBUSHOME/bin/nco_sql -user user_name -server server_name < path_to_extracted_file/apm/sql/apm.sql
```

where:

*user\_name* is the username for logging into the ObjectServer

*server\_name* is the name of your ObjectServer for OMNIBus

*path\_to\_extracted\_file* is the path where the `ibm-cem-noi-config.zip` file was extracted

- Go to the host where your Message Bus Probe is installed. Extract the downloaded `ibm-cem-noi-config.zip` file and run the following command to start the probe:

```
$OMNIHOME/probes/nco_p_message_bus -rulesfile \  
path_to_extracted_files/apm/rules/apm.rules -propsfile \  
path_to_extracted_files/apm/props/apm.props -transportfile \  
path_to_extracted_files/apm/transport/httpTransport.properties -messagelog stdout \  
-messagelevel debug
```

**Note:** The probe listens on port 10000 by default. You can change the port number by editing the `httpTransport.properties` file that you downloaded as part of step “5” on page 718 and changing the `serverPort` value.

- Go to Cloud App Management and open the integration you saved in the earlier steps (see tile under the name you provided). Enter the webhook URL for your Message Bus Probe in the **Enter Netcool/OMNIBus webhook URL** field.  
Use the following format: `http://Netcool/OMNIBus_host:10000`, where *Netcool/OMNIBus\_host* is the name of the host where your Message Bus Probe is installed.  
The probe listens on port 10000 by default. If you changed the default port setting in step “8” on page 719, then ensure you use the port number you set there.
- Click **Save** to save the integration in Cloud App Management.
- Optional: You can secure the integration between the Message Bus Probe and Cloud App Management by setting up authentication and encryption using the following process:
  - Go to the host where your Message Bus Probe is installed.
  - Create a new directory called `keystores` in `$OMNIHOME` and change to the location, for example:

```
mkdir $OMNIHOME/keystores  
cd $OMNIHOME/keystores
```

- Use the **keytool** utility to configure your authentication, for example:

```
$NCHOME/platform/arch/jre_1.6.7/jre/bin/keytool \  
-genkey -alias localhost -keystore probe.jks -storetype JKS -keyalg rsa -dname \  
"CN=localhost, OU=Hybrid Cloud, O=IBM, L=London, S=London, C=UK" -storepass \  
probepw -keypass probepw
```

Change the parameters as required for your environment. If you do not specify the `-dname` values, you are prompted for each:

- CN: CommonName
- OU: OrganizationalUnit
- O: Organization
- L: Locality
- S: StateOrProvinceName
- C: CountryName

**Important:** The CN value must be set to the host where your Message Bus Probe is installed. This is the same host name value you provided in the webhook URL as part of step “9” on page 719.

- Create a copy of the `httpTransport.properties` file you downloaded as part of step “5” on page 718, and rename the copy to `httpsWebhookTransport.properties`.
- Edit the `httpsWebhookTransport.properties` file and configure the following properties:

```
serverPort=https:10000  
keyStore=full_path_to_$OMNIHOME/keystores/probe.jks  
keyStorePassword=probepw
```

```
username=username_without_quotes
password=password_without_quotes
```

- f) Edit the `apm.props` file you downloaded as part of step “5” on page 718, and add the following property:

```
TransportFile : 'path_to_file/httpsWebhookTransport.properties'
```

- g) Start the probe as described in “8” on page 719, but set the `-transportfile` value to the `httpsWebhookTransport.properties` file you configured: `-transportfile / path_to_file/httpsWebhookTransport.properties`.
- h) Go to Cloud App Management and open the integration you saved in the earlier steps (see tile under the name you provided). Edit the following values and click **Save**:
- Go to the **Enter Netcool/OMNIbus webhook URL** field and change `http` to `https` in the webhook URL.
  - Go to the **Enter your credentials** section and enter the user name and password you provided in the `httpsWebhookTransport.properties` file.
- Note:** For more information about enabling a secure webhook integration with the Message Bus Probe, see [Probes and probe integrations](#) and [Message Bus Probe for Webhook Integrations Helm Chart](#) in the IBM Knowledge Center.
12. To send notifications from Cloud App Management to Netcool/OMNIbus, ensure that **Enable integration** is set to **On** in Cloud App Management.
13. Set up an incident policy where you set your Netcool/OMNIbus integration as a recipient of notifications. For more information, see [“Managing incident policies”](#) on page 742.

## Results

Once this integration and policy are activated, Cloud App Management will only send new incidents to Netcool/OMNIbus with the details of the first event. Cloud App Management will not send or create new rows in Netcool/OMNIbus for existing incidents in Cloud App Management, even if those Cloud App Management incidents are subsequently updated with new events on the Cloud App Management incident feed. As a result, you may see incidents on the incident feed with a recent timestamp, but no corresponding incident on the Netcool/OMNIbus event viewer.

## Sending incident details to Slack channels

You can send notifications to Slack channels. Slack is a cloud-based team collaboration tool that facilitates real-time messaging and file sharing.

### Before you begin

If you want to send notifications to your Slack channels from IBM Cloud App Management in an IBM Cloud Private environment, you must configure an incoming WebHook URL within your Slack service. The WebHook URL provided by Slack is required for the integration steps later in this section. Complete the following steps to create the WebHook URL:

1. From your Slack channel click the icon for **Channel Settings** > **Add apps** and search for "incoming-webhook".
2. Click **Add configuration**.
3. Select the channel that you want to post to.
4. Click **Add Incoming WebHooks integration**.
5. Copy the URL in the **WebHook URL** field and paste it in the field provided on the IBM Cloud App Management Slack integration page.
6. Click **Save Settings**.

## About this task

### Note:

IBM is providing a link to this third party channel as a convenience to you.

IBM does not control the processing or security of that third party channel and is not responsible for that processing or security.

You are responsible for determining whether or not the third party channel meets your requirements in terms of processing, security and privacy.

To set up a Slack channel integration:

### Procedure

1. Click **Integrations** on the **Administration** page.
2. Click the **Outgoing** tab, and click **Configure an integration**.
3. Go to the **Slack** tile and click **Configure**.
4. When adding a team for the first time you must complete two additional steps. These steps also apply when you select **Change teams**.
  - a) Enter your team's Slack domain and click **Continue**.
  - b) Enter your Slack credentials. Your email address and Slack password are required to allow the Slack App to access the available channels so that you can add them.
5. Select a Team and Channel to post to. You can have multiple teams and switch between them. If you have more than one team the list of channels displayed is for the currently selected team.
6. Click **Authorize**. The Slack App is added to the team and the channel appears in the **Slack integration** list in IBM Cloud App Management.
7. On the **Integrations** page, ensure that you set **Enablement** to **On** for **Slack** to allow the Slack channel to receive notifications from IBM Cloud App Management.
8. To send notifications about incidents to Slack channels, set up Incident policies as described in [“Managing incident policies” on page 742](#).

### What to do next

- Notifications from IBM Cloud App Management to your Slack channels include information about the incident such as incident state, priority, description, and links to view the incident details in the Cloud Event Management UI.
- To delete an existing Slack integration, on the **Integrations** page click the actions menu on the appropriate Slack integration tile and select **Delete**.

## Sending incident details using outgoing webhooks

You can use outgoing webhooks to connect to third party applications and services, and send them notifications from IBM Cloud App Management. The outgoing webhooks integration is only available in IBM Cloud App Management, Advanced.

### About this task

After setting up an outgoing webhook integration, use incident policies to post incident details to third party applications and services via the webhook.

### Note:

IBM is providing a link to this third party channel as a convenience to you.

IBM does not control the processing or security of that third party channel and is not responsible for that processing or security.

You are responsible for determining whether or not the third party channel meets your requirements in terms of processing, security and privacy.

## Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click the **Outgoing** tab, and click **Configure an integration**.
3. Go to the **Webhook** tile and click **Configure**.
4. In the **Create an outgoing webhook** window, provide the details for the integration:
  - a) In the **Name** field, enter a name for the outgoing webhook integration.
  - b) In the **Outgoing webhook URL** field, provide the webhook from the third party application or service. This is the URL that is used to send information to that application or service from IBM Cloud App Management.
  - c) Optional: In the **Basic authentication username** and **Basic authentication password** fields, enter the user name and password if basic authentication is required. If no authentication is needed, then leave blank.
  - d) Set **Enable integration** to **On**.
  - e) Click **Save**. The outgoing webhook integration is added to the **Outgoing webhook integration** list.
5. On the **Integrations** page, ensure you set **Enablement** to **On** for **Outgoing webhook** to send notifications from IBM Cloud App Management using webhooks.
6. To send notifications about incidents using outgoing webhooks, set up Incident policies as described in [“Managing incident policies”](#) on page 742.

## What to do next

- Using outgoing webhooks, the notifications from IBM Cloud App Management are sent as JSON files to other applications and services, and include information about the incident such as incident state, priority, description, and links to view the incident details in the Cloud Event Management UI.
- To edit the properties of the outgoing webhook integration, on the **Integrations** page click the actions menu on the appropriate webhook integration tile and select **Edit**.
- To delete an existing outgoing webhook integration, on the **Integrations** page click the actions menu on the appropriate webhook integration tile and select **Delete**.

## Sending incident details to Microsoft Teams

Configure this outgoing integration to send incident information to a Microsoft Teams channel from notifications when added as a recipient to an incident policy, and that incident policy matches an incident. The Microsoft Teams integration is only available in IBM Cloud App Management, Advanced.

### About this task

#### Note:

IBM is providing a link to this third party channel as a convenience to you.

IBM does not control the processing or security of that third party channel and is not responsible for that processing or security.

You are responsible for determining whether or not the third party channel meets your requirements in terms of processing, security and privacy.

To set up a Microsoft Teams integration:

## Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click the **Outgoing** tab, and click **Configure an integration**.
3. Go to the **Microsoft Teams** tile and click **Configure**.
4. In the **Create the outgoing integration with Microsoft Teams** window, provide the details for the integration:

- a) In the **Name** field, enter a name for the outgoing Microsoft Teams integration.
  - b) In the **Enter a Microsoft Teams webhook URL** field, provide the webhook from Microsoft Teams. This is the URL that is used to send information to Microsoft Teams from IBM Cloud App Management. To obtain the webhook URL:
    - 1) Open Microsoft Teams and click the **Store** icon on the sidebar.
    - 2) Type `incoming` webhook in the search box and select the **Incoming Webhook** connector.
    - 3) Follow the on-screen prompts to select a team, channel, and name for IBM Cloud App Management.
    - 4) After creating the webhook, you are given a URL to copy and paste here.
5. Set **Enable integration** to **On**.
  6. Click **Save**. The Microsoft Teams integration tile is added to the outgoing integrations page.
  7. On the **Integrations** page, ensure that you set **Enablement** to **On** for **Microsoft Teams** to allow Microsoft Teams to receive notifications from IBM Cloud App Management.
  8. To send notifications about incidents to Microsoft Teams channels, set up Incident policies as described in [“Managing incident policies” on page 742](#).

### What to do next

- Notifications from IBM Cloud App Management to Microsoft Teams include information about the incident such as incident state, priority, description, and links to view the incident details in the IBM Cloud App Management UI.
- To delete an existing Microsoft Teams integration, on the **Integrations** page click the actions menu on the Microsoft Teams tile and select **Delete**.

## Sending incident details to ServiceNow

Configure this outgoing integration to send incident information to a ServiceNow environment from the IBM Cloud App Management when added as a recipient to an incident policy, and that incident policy matches an incident. The ServiceNow is only available in IBM Cloud App Management, Advanced.

### About this task

#### Note:

IBM is providing a link to this third party channel as a convenience to you.

IBM does not control the processing or security of that third party channel and is not responsible for that processing or security.

You are responsible for determining whether or not the third party channel meets your requirements in terms of processing, security and privacy.

To set up a ServiceNow integration:

### Procedure

1. Click **Integrations** on the Cloud Event Management **Administration** page.
2. Click the **Outgoing** tab, and click **Configure an integration**.
3. Go to the **ServiceNow** tile and click **Configure**.
4. In the **Create the outgoing integration with ServiceNow** window, provide the details for the integration:
  - a) In the **Name** field, enter a name for the outgoing ServiceNow integration.
  - b) In the **Enter a ServiceNow webhook URL** field, provide the webhook from ServiceNow. This is the URL that is used to send information to ServiceNow from IBM Cloud App Management. Use the following format for the URL:

```
https://<instanceName>.service-now.com/api/now/table/incident
```

Where <instanceName> is your unique ServiceNow instance name.

- c) Enter your ServiceNow user name and password in the fields provided.
5. Set **Enable integration to On**.
6. Click **Save**. The ServiceNow integration tile is added to the outgoing integrations page.
7. On the **Integrations** page, ensure that you set **Enablement to On** for **ServiceNow** to allow ServiceNow to receive notifications from IBM Cloud App Management.
8. To send notifications about incidents to a ServiceNow environment, set up Incident policies as described in [“Managing incident policies” on page 742](#).

### What to do next

- Notifications from IBM Cloud App Management to ServiceNow include information about the incident such as incident state, priority, description, and links to view the incident details in the IBM Cloud App Management UI.
- To delete an existing ServiceNow integration, on the **Integrations** page click the actions menu on the ServiceNow tile and select **Delete**.

## Sending incident details to GitHub

Configure this outgoing integration to send incident information to a GitHub repository as a GitHub issue when added as a recipient to an incident policy, and that incident policy matches an incident. The GitHub integration is only available in IBM Cloud App Management, Advanced.

### About this task

#### Note:

IBM is providing a link to this third party channel as a convenience to you.

IBM does not control the processing or security of that third party channel and is not responsible for that processing or security.

You are responsible for determining whether or not the third party channel meets your requirements in terms of processing, security and privacy.

To set up a GitHub integration:

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click the **Outgoing** tab, and click **Configure an integration**.
3. Go to the **GitHub** tile and click **Configure**.
4. In the **Create the outgoing integration with GitHub** window, provide the details for the integration:
  - a) In the **Name** field, enter a name for the outgoing GitHub integration.
  - b) In the **Enter a GitHub webhook URL** field, provide the webhook from GitHub. This is the URL that is used to send information to GitHub from IBM Cloud App Management.

Use the following format for the URL:

```
https://<hostname>/api/v3/repos/<organizationName>/<repoName>/issues
```

- c) Enter your GitHub user name and API Token in the fields provided.

To create a GitHub API token, go to GitHub and click **Settings > Developer settings > Personal access tokens** and **Generate new token** with the **repo** scope selected.
5. Set **Enable integration to On**.
6. Click **Save**. The GitHub integration tile is added to the outgoing integrations page.



7. On the **Integrations** page, ensure you set **Enablement to On** for **GitHub** to allow GitHub to receive notifications from IBM Cloud App Management.
8. To send notifications about incidents to GitHub, set up Incident policies as described in [“Managing incident policies”](#) on page 742.

### What to do next

- Notifications from IBM Cloud App Management to GitHub include information about the incident such as incident state, priority, description, and links to view the incident details in the IBM Cloud App Management UI.
- To delete an existing GitHub integration, on the **Integrations** page click the actions menu on the GitHub tile and select **Delete**.

## Sending incident details to Stride

Configure this outgoing integration to send incident information to a Stride channel from notifications when added as a recipient to an incident policy, and that incident policy matches an incident. The Stride integration is only available in IBM Cloud App Management, Advanced.

### Before you begin

You will need a Stride account, team, and channel to create an integration with Cloud Event Management.

### About this task

#### Note:

IBM is providing a link to this third party channel as a convenience to you.

IBM does not control the processing or security of that third party channel and is not responsible for that processing or security.

You are responsible for determining whether or not the third party channel meets your requirements in terms of processing, security and privacy.

To set up a Stride integration:

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click the **Outgoing** tab, and click **Configure an integration**.
3. Go to the **Stride** tile and click **Configure**.
4. Enter a name for the outbound integration.
5. Complete the following steps to obtain a Stride webhook URL and access token:
  - a) Go to the Stride channel that you want CEM to notify and then click **Apps > Add an App > Add custom app > API Tokens**.
  - b) Specify a token name and click **Create**.
  - c) Copy and paste the **API URL** and **access token** into the fields provided in steps 2 and 3.
6. Set **Enable integration to On**.
7. Click **Save**. The Stride integration tile is added to the outgoing integrations page.
8. On the **Integrations** page, ensure that you set **Enablement to On** for **Stride** to allow Stride to receive notifications from IBM Cloud App Management.
9. To send notifications about incidents to Stride channels, set up Incident policies as described in [“Managing incident policies”](#) on page 742.

## What to do next

- Notifications from IBM Cloud App Management to Stride include information about the incident such as incident state, priority, description, and links to view the incident details in the IBM Cloud App Management UI.
- To delete an existing Stride integration, on the **Integrations** page click the actions menu on the Stride tile and select **Delete**.

## Sending incident details to Watson Workspace

Configure this outgoing integration to send incident information to Watson Workspace from notifications when added as a recipient to an incident policy, and that incident policy matches an incident. The Watson workspace integration is only available in IBM Cloud App Management, Advanced.

### Procedure

1. Click **Integrations** on the IBM Cloud App Management **Administration** page.
2. Click the **Outgoing** tab, and click **Configure an integration**.
3. Go to the **Watson Workspace** tile and click **Configure**.
4. In the **Create the outgoing integration with Watson Workspace** window, provide the details for the integration:
  - a) In the **Name** field, enter a name for the outgoing Watson Workspace integration.
  - b) Click **Add to Watson Workspace**.
  - c) Select the appropriate workspace and click **Add App**.
  - d) On the message confirming that IBM Cloud App Management was added to the workspace, click **Go to space**.
  - e) In the **Spaces** pane on the left of Watson Workspace, select the space CEM should notify and then click the more options icon **>** **Copy space link**.
  - f) Paste the link into the field provided in step 3 of the integration window.
5. Set **Enable integration** to **On**.
6. Click **Save**. The Watson Workspace integration tile is added to the outgoing integrations page.
7. On the **Integrations** page, ensure that you set **Enablement** to **On** for **Watson Workspace** to allow Watson Workspace to receive notifications from IBM Cloud App Management.
8. To send notifications about incidents to Watson Workspace, set up Incident policies as described in [“Managing incident policies” on page 742](#).

## What to do next

- Notifications from IBM Cloud App Management to Watson Workspace include information about the incident such as incident state, priority, description, and links to view the incident details in the IBM Cloud App Management UI.
- To delete an existing Watson Workspace integration, on the **Integrations** page click the actions menu on the Watson Workspace tile and select **Delete**.

---

# Chapter 19. Administering

Administering IBM Cloud App Management involves tasks that range from setting up users and groups to managing incident policies.

## Setting up event policies

You can set up event policies to handle a set of events in a specified way. You can determine what events you want the policy to apply to, and select one or more actions to take on those events. For example, you could choose actions to suppress events or to assign runbooks to events.

### About this task

To create an event policy:

### Procedure

1. Click **Policies** on the IBM Cloud App Management **Administration** page.
2. Click **Create event policy**.
3. Enter a name and a description for the policy in **Details**.
4. Specify what events you want the policy to apply to in **Events**. You can specify to have all events considered for the policy actions by clicking **All events**, or you can configure what conditions the events have to meet before the actions are applied to them by clicking **Specify conditions**.

**Tip:** When selecting **Specify conditions**, you can join multiple conditions using the AND and OR operators. You can also use the example conditions provided by clicking **Use example**. To view the examples, expand **Information and examples > Show examples**. In addition, you can select from a list of predefined conditions to use by clicking **Add predefined condition**.

5. Optional: When selecting **Specify conditions**, you can check to see how many events would have matched the conditions you set. Go to the end of the **Events** section, select the number of days between 1 and 30, and click **Test**. The result shows how many events would have matched the policy conditions.

Click **Show results** to view a list of all the events that would have matched the conditions in the set time. Click **New test** to change the time frame for testing, or if you changed conditions and want to check again for matching events.

**Note:** If your event policy enriches fields used by the conditions of your policy, you might not find any matching events after the policy is enabled and applied.

6. Go to the **Action** section, and set what actions you want the policy to take against the events.

- **Enrich:** Change existing event information or add new information to the event.

**Tip:** Event enrichment can be used to correlate events into incidents. See example scenario later.


- **Suppress:** Set whether all events specified in the previous step are suppressed, or only in case a specified number of them occur within a set time frame. Suppressing events stop them from forming an incident or becoming part of existing incidents.

- **Assign runbooks:** Specify which runbooks are available to run against the specified events.

**Important:** To assign runbooks, ensure you have runbooks that are published to make them available to event policies. For more information, see [“Managing runbooks” on page 764](#).

Runbooks can be run manually or automatically. When assigning manual runbooks, you can set whether you want parameter values for the runbook to be taken from the event, entered manually, or specified at runtime. Automatic runbooks contain only automated steps and you must select the **Automatically run this runbook** check box when assigning the runbook to events in the event policy. Automatic runbooks can only take parameter values from the events or if you provide them


when setting up the policy. Ensure you select **From event** or **Manual input** for the parameter settings, and set the appropriate values.

- **Detect flapping:** Mark events that close and reopen rapidly as flapping events. Flapping events point to recurring problems, and are noted in the incident **Events** tab with the  **Event is flapping** icon to highlight the condition. When an incident contains flapping events, it cannot be resolved automatically until the events stop flapping, even if all other events that form part of the incident are cleared. This is to ensure that the root cause of any flapping event is investigated and rectified before the incident can be declared as resolved. If a user tries to manually set an incident with flapping events to resolved, they are warned that flapping events might cause the incident to reopen.

**Tip:** Cloud Event Management provides a built-in event policy called **Global flapping detection** to identify flapping events. The policy detects events that clear and reopen 4 or more times in an hour, and marks them as flapping. If these events stop changing states for more than 30 minutes, they are no longer considered to be flapping. This policy applies to all events and is enabled by default. To view this policy, go to the Cloud Event Management **Administration** page, click **Policies**, and ensure you are on the **Event policy** tab. Look for **Global flapping detection** in the list of event policies. You can use the built-in flapping policy to detect flapping events, or set up your own as described in the example scenario later.

- **Forward events:** The event will not create an incident in Cloud Event Management, but instead will be forwarded to the specified integration. Note, when event forwarding is enabled, **Suppress**, **Assign runbook**, and **Detect flapping** actions will not be applied to the event.

See the scenarios later for examples of using these actions to set up different policies against events.

7. Set **Enable** to **On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
8. Click **Save**. You are returned to the list of event policies.
9. You can set the order in which your policies are applied. Using the  **Menu overflow**, you can move any selected policy up or down the list, or move it to the top or bottom. The numbering determines the ranking with 1 being the highest priority.

## Events and incidents

Use the Cloud APM console to set up real-time incident management for your services, applications, and infrastructure.

### Events, incidents, and correlation

Events indicate that something has happened on an application, service, or another monitored object.

Related events are correlated into an incident based on attribute values that match. This means that events such as alerts or notifications from monitoring tools are considered to be part of the same incident if they have the same information set in a specific attribute.

Events are *deduplicated*, meaning that if the same event occurs multiple times, not all are listed in the incident they are correlated into. Instead, the count for the same event is updated to show how many times it has occurred. By default, the event severity, summary, state, and type fields are updated on deduplication to always have the latest values.

Events that have matching values in one of the following attributes are correlated into an incident. The attributes are checked in the following order:

- Cluster (event.resource.cluster)
- Application (event.resource.application)
- Hostname (event.resource.hostname)
- IP address (event.resource.ipaddress)
- Resource source ID (event.resource.sourceId)
- Service (event.resource.service)

- Resource name (event.resource.name)

The value in the Cluster attributes is checked first. If the Cluster values are the same, then the events are correlated into an incident. If the Cluster values are different, then no correlation takes place.

If there is no value set in the Cluster attribute for any of the events, the process moves on to the Application attribute and checks whether the values match there. If the Application values are the same, then the events are correlated into an incident. If the Application values are different, then no correlation takes place.

If there is no value set in the Application attribute for any of the events, the process moves on to check the Hostname values. The process continues through the list and checks the attribute values until one of them match across events, and none of the previous attributes have values. When the same event attribute has the same value for more than one event, those events are correlated into an incident.

The following example shows attribute values for several events. The process uses the values to check whether any of the events can be correlated into incidents.

*Table 95. Example: Event attribute values for correlation*

Attribute	Event A	Event B	Event C	Event D	Event F	Event G	Event H
event.resource .cluster		cluster A		cluster A			cluster B
event.resource .application		payroll		billing	payroll		payroll
event.resource .hostname	hostA1			hostA1	hostA1		hostA1
event.resource .ipaddress	1.1.1.1	1.1.1.1		2.2.2.2			2.2.2.2
event.resource .sourceId	ABC1						ABC1
event.resource .service	web	web			web		web
event.resource .name	databaseA		databaseA			databaseA	databaseA

Using the process described previously, only the following events are correlated in this example:

- Events B and D are correlated into an incident based on identical Cluster values. The correlation is based only on Cluster values as the top-level attribute, and no other attribute values are considered afterward as part of correlation.
- Events C and G are correlated into an incident based on identical Resource name values, and because no other higher-level attributes have any values set.

The other events do not correlate into any incident due to the following:

- Event A does not correlate with any other event despite its Hostname value matching Event D's Hostname value. The previous attribute values are not set in Event A, while they are set in Event D.
- Event F does not correlate with any other event despite its Application value matching Events B's Application value. Event B has the Cluster attribute set, while Event F has no value set for Cluster.

- Event H does not correlate with any other event as it has a value set in all attributes, and the Cluster value is unique to the event.

Investigating and resolving the incident solves the underlying problem that caused the events to form the incident, and restores service.

### **Event severity and incident priority**

To help prioritize and manage problems efficiently, events have severity levels that are used in ranking the importance of incidents.

Events can have the following severity levels, depending on how serious the problem the event relates to is:

- Indeterminate
- Information
- Warning
- Minor
- Major
- Critical

The severity levels are based on the alert information from the event source. By default, event severity determines how incidents are ranked in importance, setting the priority for the incident.

Incident priority ranges from 1 to 5, with 1 being the highest priority. The priority of the incident is based on the severity of the events that make up the incident, with the highest severity event determining the overall priority of the incident. The incident priority is set as follows:

- Priority 1: if an incident contains critical severity level events.
- Priority 2: if an incident contains major severity level events.
- Priority 3: if an incident contains minor severity level events.
- Priority 4: if an incident contains warning severity level events.
- Priority 5: if an incident contains information or indeterminate severity level events.

**Important:** By default, built-in incident policies set the priority of incidents as described here. If you modify or remove the built-in incident policies that set priority, or add new policies, then the described behavior changes and the incident priorities are set based on your adjustments to the policies. For more information see [“Managing incident policies” on page 742](#).

### **Event sources**

You can configure an integration with tools and systems from which you want to receive events. You can configure the integration by defining one or more event sources.

Events can be obtained from various sources. For more information, see [“Configuring incoming event sources” on page 692](#).

### **Policies**

You can create policies to take action against incidents. The policy actions help you manage problems more efficiently. Incident policies act on incidents, such as to assign them to specified groups automatically, notify users automatically, or escalate ones that have no investigation in progress after a specified period of time.

### **Runbooks**

The information from incidents helps operations teams respond to service problems. You can use runbooks to improve efficiency by capturing knowledge of similar incidents over time, and building guidance and automation for resolving them. Runbooks provide structured manual and automated steps to help solve the underlying problems that are described in the incidents, so you can restore service fast.

## Users and groups

You can invite your team members to Cloud App Management and organize them into groups. Incidents can be routed to the groups with the expertise to resolve them, for example, database administrators. The groups can assign the incidents to the appropriate individuals.

You can also create policies for your teams that determine when they receive notifications about incidents and specify methods for how they are notified of those incidents.

## Example: Changing event information through enrichment

You can change event data using the enrich action in the event policy. Changing specific information provided by events can help address issues more efficiently in some scenarios.

### About this task

For example, you might not always have control over the severity level of the events generated by the monitoring tool. In some cases you might want to change the actual severity of the problem.

You could have a monitoring tool that generates a warning event when high CPU usage is detected. The event has a severity level of Major. You cannot change how the monitoring tool sets the severity. However, you might want the severity for such issues to be increased to Critical to ensure that the underlying issue receives the right attention before it causes other problems. Using the enrich action, you can set up an event policy that changes the severity of such events to Critical.

To set up this policy:

### Procedure

1. Click **Policies** on the IBM Cloud App Management **Administration** page.
2. Click **Create event policy**.
3. Go to **Details** and enter a name in **Policy name**, for example, Change severity for high CPU usage events. You can also add an explanation of the policy in **Description** to help you and others understand the purpose of the policy, for example, Change severity level for high CPU events to critical to ensure they receive prompt attention.
4. Click **Specify conditions in Events**, and set the following conditions:
  - a) Set **Condition 1** as follows: select **Sender type** from the list of attributes, select **is** from the list of operators, and enter the name of the monitoring tool in the field, for example, Datadog.  
**Tip:** If you have more than one instance set up for the same monitoring tool, and you only want to enrich events from one of them, you can use the **Sender display name** instead. The **Sender display name** value is mapped to the name provided in the IBM Cloud App Management UI when setting up the integration with the event source.  
**Note:** This is an example. The attribute values depend on your event source. When creating similar policies, check the values from your events to ensure you set the correct value.
  - b) Ensure you have **AND** set and click **Add condition**.
  - c) Set **Condition 2** as follows: select **Event Type** from the list of attributes, select **is** from the list of operators, and enter CPU\_HIGH in the field.  
**Note:** This is an example. The attribute values depend on your event source. When creating similar policies, check the values from your events to ensure you set the correct value.
  - d) Ensure you have **AND** set again and click **Add condition**.
  - e) Set **Condition 3** as follows: select **Severity** from the list of attributes, select **is** from the list of operators, and select **Major**.
5. Optional: When selecting **Specify conditions**, you can check to see how many events would have matched the conditions you set. Go to the end of the **Events** section, select the number of days between 1 and 30, and click **Test**. The result shows how many events would have matched the policy conditions.

Click **Show results** to view a list of all the events that would have matched the conditions in the set time. Click **New test** to change the time frame for testing, or if you changed conditions and want to check again for matching events.

**Note:** If your event policy enriches fields used by the conditions of your policy, you might not find any matching events after the policy is enabled and applied.

6. Select the **Enrich** check box in **Action**, and expand the section.
7. Select **Severity** from the list of attributes, and then **Critical** from the **Select severity** list.
8. Set **Enable** to **On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
9. Click **Save**.

## Results

When events match the set conditions, the severity value is changed to Critical.

## Example: Adding to event information through enrichment

You can add to event data using the enrich action in the event policy. Adding information to specific events can help inform your teams about the problem more accurately.

### About this task

For example, you might have a monitoring tool that sends a short summary included in the event data. The summary might be a basic note about the problem which does not carry enough detail to make it clear what the issue is. Using the enrich action, you can add more detail to the summary, making it more helpful in understanding the issue at a glance.

A possible example is when critical warnings about high bandwidth utilization only include a short summary stating "band util critical". Using the enrich action, you can set up an event policy that adds information to the summary to make it more meaningful.

### Procedure

1. Click **Policies** on the IBM Cloud App Management **Administration** page.
2. Click **Create event policy**.
3. Go to **Details** and enter a name in **Policy name**, for example, Bandwidth warnings: Make summary more informative. You can also add an explanation of the policy in **Description** to help you and others understand the purpose of the policy, for example, Update summary for high bandwidth utilization events to make them more meaningful. Apply to critical or more severe warnings.
4. Click **Specify conditions** in **Events**, and set the following conditions:
  - a) Set **Condition 1** as follows: select **Event Type** from the list of attributes, select **is** from the list of operators, and enter the identifier for the type of event in the field: BAND\_UTIL.  
**Note:** This is an example. The attribute values depend on your event source. When creating similar policies, check the values from your events to ensure you set the correct value.
  - b) Ensure you have **AND** set and click **Add condition**.
  - c) Set **Condition 2** as follows: select **Severity** from the list of attributes, select **Is** from the list of operators, and select **Critical**.
5. Optional: When selecting **Specify conditions**, you can check to see how many events would have matched the conditions you set. Go to the end of the **Events** section, select the number of days between 1 and 30, and click **Test**. The result shows how many events would have matched the policy conditions.  
Click **Show results** to view a list of all the events that would have matched the conditions in the set time. Click **New test** to change the time frame for testing, or if you changed conditions and want to check again for matching events.



**Note:** If your event policy enriches fields used by the conditions of your policy, you might not find any matching events after the policy is enabled and applied.

6. Select the **Enrich** check box in **Action**, and expand the section.
7. Select **Summary** from the list of attributes, and enter the following text in the field to update the summary with: **Bandwidth utilization for the interface is critically high. Application response times may be affected.** Ensure you have **Append to field** selected.
8. Set **Enable to On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
9. Click **Save**.

## Results

When events match the set conditions, the summary for such events is updated with the text provided. In this example, the text is added after the existing summary description. You can also select to add it before the description that arrives with the event, or overwrite the summary entirely with the description you specify.

## Example: Correlating events through enrichment

You can correlate events using the enrich action in the event policy. Correlation through enrichment can ensure events relating to the same incident are grouped together correctly.

### About this task

In some cases your events might be missing information that could be used to correlate them with other related events, making them form part of the same incident.

For example, you might have a hybrid application called **My Hybrid App**, with the front-end client side hosted on a cloud platform, and the back-end servers hosted on premises. The events from the front end have the application information set in the monitoring tool. However, the back end servers do not have the same application information set. This means that the events from the front end and back end cannot be correlated, even though they all relate to the same application. The lack of information on the back-end servers is preventing all related events to be correlated into the incident affecting the same application. This could lead to difficulty in understanding the problem the incident is about, as you might have warnings from the front end that are caused by back end problems that are not flagged in the same incident.

To have both front-end and back-end events correlated, you can use enrichment to add the missing application information to the events coming from the back-end monitoring tools. Together with the same application information already set in the front-end events, all events related to the application can be then correlated into an incident.

**Note:** This scenario assumes that events from the back-end servers can be identified in a unique way. In this example, that identification is based on host names.

### Procedure

1. Click **Policies** on the Cloud Event Management **Administration** page.
2. Click **Create event policy**.
3. Go to **Details** and enter a name in **Policy name**, for example, **Set application field for event correlation**. You can also add an explanation of the policy in **Description** to help you and others understand the purpose of the policy, for example, **Set application information for events from back-end servers to enable correlation with front-end events relating to same application**.
4. Click **Specify conditions** in **Events**, and set the following conditions:
  - a) Set **Condition 1** as follows: select **Hostname** from the list of attributes, select **is** from the list of operators, and enter the host name of the first back-end server in the field, for example `abc.div.org.com`.

- b) Ensure you have **OR** set and click **Add condition**.
  - c) Set **Condition 2** as follows: select **Hostname** from the list of attributes, select **is** from the list of operators, and enter the host name of the second back-end server in the field, for example `def.div.org.com`.  
  
**Remember:** This scenario assumes the back-end servers can be identified by their host names. The custom conditions for the event are set to identify events coming from either of the two hosts for the back-end servers.
5. Optional: When selecting **Specify conditions**, you can check to see how many events would have matched the conditions you set. Go to the end of the **Events** section, select the number of days between 1 and 30, and click **Test**. The result shows how many events would have matched the policy conditions.  
Click **Show results** to view a list of all the events that would have matched the conditions in the set time. Click **New test** to change the time frame for testing, or if you changed conditions and want to check again for matching events.  
  
**Note:** If your event policy enriches fields used by the conditions of your policy, you might not find any matching events after the policy is enabled and applied.
  6. Select the **Enrich** check box in **Action**, and expand the section.
  7. Select **Application** from the list of attributes, and enter the same application information in the field as set for the front-end events, for example: **My Hybrid App**. Ensure you have **Overwrite field** selected.
  8. Set **Enable** to **On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
  9. Click **Save**.

## Results

When events match the set conditions, the **Application** value for such events is updated to include the name provided. In this example, the value **My Hybrid App** is set for the events coming from the specified back-end servers. This enables Cloud Event Management to correlate the back end events with the front end events that relate to the same application, creating a single incident for all events relating to the same problem on the application.

## Creating lookup tables

Lookup tables are used to enable the fast and easy lookup of static data. You can use lookup tables to enrich events by correlating attributes in the events with corresponding attributes in the lookup table.

### About this task

Import the contents of a lookup table in CSV format. A basic example is a CSV file containing application names and a summary update. This data might be used to add summary information to events. In the enrichment example that follows this topic, the following CSV file is imported to create the lookup table.

```
applicationname,summaryupdate
HR,-HR Application Affected
Human Resources,-HR Application Affected
Payroll Application,-Payroll Application Affected
```

### Procedure

1. Click **Lookup tables**.
2. Click **Lookup tables > New lookup table +**.
3. Enter a name and a description for the lookup table in **Details**.
4. Click **Import from CSV** and browse to your CSV file to upload the contents to Cloud App Management.
5. Click **Save**.

For more information, see [“Example: Enriching event information using lookup tables” on page 735](#).

## Example: Enriching event information using lookup tables

Lookup tables use information in your events to determine how to add other fields from external data sources such as CSV files. Event policies can contain multiple lookup tables. Some event fields have more options than others. For instance, you can only replace the attribute *Hostname* while you can prepend, append, and replace the attribute *Summary*.

### About this task

A basic example would be using the lookup table that we created in “Creating lookup tables” on page 734 with application names and a summary update to add summary information to an event. You might have a monitoring tool that sends event data about the applications, but it's not immediately clear which application is affected. Using the enrich action and lookup capability, you can add more detail to the summary, making it more helpful in understanding the issue at a glance.

Lets examine how this lookup table is applied in the following policy example. In Figure 1 the value in the event of the attribute **Application** (seen in Figure 4) is compared with the value in the column **applicationname** in Figure 1. If a match is found, in this case in row 3 (highlighted in red), then the value in the **summaryupdate** column of row 3 will be appended to the Event summary as shown in Figure 2.

Define lookup criteria

Using table: **Application Lookup Table**

Enrich [Summary] from column: **summaryupdate**

Where event attribute **Application** matches column **applicationname**

applicationname	summaryupdate
HR	-HR Application Affected
Human Resources	-HR Application Affected
Payroll Application	-Payroll Applications Affected

Cancel Apply

Figure 5. Example lookup table criteria

When events match the conditions the attributes will be modified as specified by the lookup criteria and, in this case, appended to the field.

Figure 6. Enrich via lookup

In this example the summary information -*Payroll Application Affected* is appended to the summary description field for events related to the Payroll Application.

Event info	
Name	Description
Summary	LOOKUP TABLE Example2-Payroll Application Affected

Figure 7. Resulting enriched event

Resource affected <span style="float: right;">i</span>	
Name	Description
Application	Payroll Application

Figure 8. Resource affected

### Procedure

1. Click **Policies** on the Cloud Event Management **Administration** page.
2. Click **Create event policy**.
3. Go to **Details** and enter a name in **Policy name**. You can also add an explanation of the policy in **Description** to help you and others understand the purpose of the policy.
4. In **Events**, click **All events** or click **Specify conditions** to configure what conditions the events have to meet before the enrichment is applied to them.
5. Select the **Enrich** check box under **Action**.
6. In the first field, select the event attribute that you are enriching from the list of available attributes.
7. In the second field, click the drop-down arrow and select **lookup**.
8. Click **Select lookup criteria** and use the drop-down lists to select a value for each field displayed:

#### Using table

Select an existing lookup table from the list available. For more information, see [“Creating lookup tables”](#) on page 734.

#### Enrich [the target fieldname] from column

The column that will supply the value to enrich the event attribute when the **matches columns** row value is the same as the specified event attribute.

#### Where event attribute

The event attribute used to search the table key field (or the matches column).

### matches column

The column that will be compared with the event attribute to determine the enrichment value from the **Enrich from column**.

9. Click **Apply**.
10. Set **Enable** to **On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
11. Click **Save**.

### Results

When events match the set conditions, the event information will be enriched with the correlated values from the lookup table, as specified by the lookup criteria.

## Enriching event information by adding custom attributes

You can add custom attributes to your event data by using the enrich action in an event policy. By using custom attributes to match specific criteria, the events that are produced include more helpful information, which informs you about the item that the event is reporting on. For example, add, remove, or parameterize custom attributes for the details of an application such as payload. When the payload event comes in, because it matches certain criteria, it produces useful event information, which can be used for forwarding events or categorizing them.

### Procedure

1. Click **Policies** on the IBM Cloud App Management **Administration** page.
2. Click **New event policy +**.
3. Go to **Details** and enter a name in **Policy name**. You can also add an explanation of the policy in **Description** to help you and others understand the purpose of the policy.

You might want to edit an existing event policy instead of creating a new one. Select a previously created event policy under the **Name** column and start editing it.

4. In **Events**, click **All events** or click **Specify conditions** to configure what conditions the events must meet before you apply the enrichment to them.
5. Select the **Enrich** check box under **Action**.
6. From the first list, select **Custom Field**. Enter a name for the new custom event attribute that you are adding.
7. From the second list, complete one of the following steps:
  - a) Select **lookup** and click **Select lookup criteria** and use the lists to select a value for each field displayed. For instructions, see [“Example: Enriching event information using lookup tables”](#) on page 735.
  - b) Select **eventAttribute** and from the **Select Attribute** list, choose an event attribute that you want to set to your new custom event attribute.
  - c) Select **=** and enter your own value to the new custom event attribute.
8. Ensure that you selected the **Overwrite field**.

In this procedure, a new custom event attribute is added. You can also add a prefix to or append (add) details to an existing attribute. For instructions, see [“Example: Adding to event information through enrichment”](#) on page 732.
9. Set **Enable** to **On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
10. Click **Save**.

### Results

A new custom event attribute is added to your event policy. When events match the set conditions and the new attribute details, the event information is enriched.

## Example: Suppressing temporary high memory usage warnings

You might want to suppress certain events to stop them from forming an incident or becoming part of existing incidents. Suppressing events can prevent effort spent on temporary issues that do not present a persistent problem.

### About this task

For example, you might have brief periods of high application usage that cause the database memory consumption to grow over normal levels. Unless the high consumption levels become persistent, you can ignore such spikes. To avoid getting distracted by events warning about high memory consumption, you can set a policy to suppress such events for the database server unless more than 5 corresponding events are raised within 15 minutes.

**Important:** This is an example. Consider your setup and environment when determining the conditions for any policy.

To set up this policy:

### Procedure

1. Click **Policies** on the IBM Cloud App Management **Administration** page.
2. Click **Create event policy**.
3. Go to **Details** and enter a name in **Policy name**, for example, Suppress temporary high memory warnings. You can also add an explanation of the policy in **Description** to help you and others understand the purpose of the policy, for example, Suppress high memory usage warnings from database server unless they become persistent.
4. Click **Specify conditions** in **Events**, and set the following conditions:
  - a) Set **Condition 1** as follows: select **Hostname** from the list of attributes, select **is** from the list of operators, and enter the host name in the field.
  - b) Ensure you have **AND** set and click **Add condition**.
  - c) Set **Condition 2** as follows: select **Event Type** from the list of attributes, select **is** from the list of operators, and enter `memoryUsage` in the field.
  - d) Click **Add condition**.
  - e) Set **Condition 3** as follows: click **Add predefined condition** and select **Severity of event is Critical**.
5. Optional: When selecting **Specify conditions**, you can check to see how many events would have matched the conditions you set. Go to the end of the **Events** section, select the number of days between 1 and 30, and click **Test**. The result shows how many events would have matched the policy conditions.

Click **Show results** to view a list of all the events that would have matched the conditions in the set time. Click **New test** to change the time frame for testing, or if you changed conditions and want to check again for matching events.

**Note:** If your event policy enriches fields used by the conditions of your policy, you might not find any matching events after the policy is enabled and applied.
6. Select the **Suppress** check box in **Action**, and expand the section.
7. Click **Suppress until** and use the arrows to select a value of 5 identical events to occur within 15 minutes.

**Note:** When an event is suppressed, no further actions from other policies are used for that event.
8. Set **Enable** to **On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
9. Click **Save**.

## Example: Preventing certain events from forming incidents

You might have situations where you want to avoid receiving events altogether, and prevent them from forming an incident. You can create policies to suppress events at all times.

### About this task

For example, you might have preproduction development and test environments that are monitored by the same monitoring tools as your production environments. The development and test environments naturally have ongoing changes that might trigger various warnings and notifications all the time. These warnings and notifications are relayed to IBM Cloud App Management as events. However, these events do not require action as they are from resources that change constantly as part of the preproduction work. The monitoring tools are usually set up to avoid sending such events to the operations teams. However, if a such event information is sent to IBM Cloud App Management, for example, due to a misconfiguration, then the operations team could face unnecessary noise.

To save your operations team from being distracted by events from such environments, you can create policies that suppress events coming from the monitored resources that make up these environments.

**Important:** This is an example. Consider your setup and environment when determining the conditions for any policy. This example assumes that the preproduction environment has naming conventions for host names.

To set up this policy:

### Procedure

1. Click **Policies** on the Cloud Event Management **Administration** page.
2. Click **Create event policy**.
3. Go to **Details** and enter a name in **Policy name**, for example, Suppress events from preproduction environment. You can also add an explanation of the policy in **Description** to help you and others understand the purpose of the policy, for example, Suppress events from hosts in the preproduction environment to prevent incidents being created for development and test hosts.
4. Click **Specify conditions** in **Events**, and set the following conditions:
  - a) Set **Condition 1** as follows: select **Hostname** from the list of attributes, select **Starts with** from the list of operators, and enter dev in the field.

**Important:** This example assumes that the preproduction environment has naming conventions for host names, with all development and test host names starting with either dev or test.
  - b) Ensure you have **OR** set, and click **Add condition**.
  - c) Set **Condition 2** as follows: select **Hostname** from the list of attributes, select **Starts with** from the list of operators, and enter test in the field.
5. Optional: When selecting **Specify conditions**, you can check to see how many events would have matched the conditions you set. Go to the end of the **Events** section, select the number of days between 1 and 30, and click **Test**. The result shows how many events would have matched the policy conditions.

Click **Show results** to view a list of all the events that would have matched the conditions in the set time. Click **New test** to change the time frame for testing, or if you changed conditions and want to check again for matching events.

**Note:** If your event policy enriches fields used by the conditions of your policy, you might not find any matching events after the policy is enabled and applied.
6. Select the **Suppress** check box in **Action**, and expand the section.
7. Click **Always suppress the described event**.
8. Set **Enable** to **On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
9. Click **Save**.

## Example: Setting runbook for disk full events

Runbooks can be used to resolve events. You might want to set specific runbooks to be available for all events, or for events that match set conditions.

### About this task

For example, you might want to have a runbook take action when you receive events warning that your disk space is filling up. The runbook could delete the contents of the /tmp directory to free up space. You can create a policy to use this runbook every time events warn of your disk space becoming full.

**Note:** This example assumes you have a runbook called **Clear tmp directories**. For information about setting up runbooks, see [“Managing runbooks” on page 764](#).

To set up this policy:

### Procedure

1. Click **Policies** on the IBM Cloud App Management **Administration** page.
2. Click **Create event policy**.
3. Go to **Details** and enter a name in **Policy name**, for example, `Clear tmp directories`. You can also add an explanation of the policy in **Description** to help you and others understand the purpose of the policy, for example, `Policy to assign a runbook to events which indicate that disk is full due to /tmp filling up`.
4. Click **Specify conditions** in **Events**, and set the following conditions:
  - a) Set **Condition 1** as follows: select **Event Type** from the list of attributes, select **is** from the list of operators, and enter `Disk_Full` in the last field.
  - b) Ensure you have **AND** set and click **Add condition**.
  - c) Set **Condition 2** as follows: select **Severity** from the list of attributes, select **Is greater than or equal to** from the list of operators, and select **Major** in the last field.
  - d) Click **Add condition**.
  - e) Set **Condition 3** as follows: click **Summary** from the list of attributes, select **Contains** from the list of operators, and enter `/tmp` in the last field.
5. Optional: When selecting **Specify conditions**, you can check to see how many events would have matched the conditions you set. Go to the end of the **Events** section, select the number of days between 1 and 30, and click **Test**. The result shows how many events would have matched the policy conditions.  
Click **Show results** to view a list of all the events that would have matched the conditions in the set time. Click **New test** to change the time frame for testing, or if you changed conditions and want to check again for matching events.  
**Note:** If your event policy enriches fields used by the conditions of your policy, you might not find any matching events after the policy is enabled and applied.
6. Select the **Assign runbook** check box in **Action**, and expand the section.
7. Click **Add runbook assignments**. The **Runbooks library** window shows the available runbooks to select, together with their name, rating by users, recorded success rate, and the execution type (manual or automated). You can use the search field to search for a runbook by name.
8. Select the runbook titled **Clear tmp directories** from the list, then click **Apply**. The selected runbook is added to the **Assign runbook** twistie.
9. You can set whether you want parameter values for the runbook to be taken from the event, entered manually, or specified at runtime. Expand the **Assign runbook** twistie, and select the runbook from the list on the left. You can then edit how the runbook takes its parameter values. In this case, leave all parameters to be taken **From event**.
10. If you want to add more runbooks, expand **Assign runbook** and click **Edit runbook assignments**.
11. Set **Enable** to **On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.




12. Click **Save**.

## Example: Detecting flapping events

Flapping events clear and reopen repeatedly in a short space of time, indicating potentially recurring problems that require investigation.

### About this task

Flapping events are noted in the incident **Events** tab with the  **Event is flapping** icon to highlight the condition. When an incident contains flapping events, it cannot be resolved automatically until the events stop flapping, even if all other events that form part of the incident are cleared. This is to ensure that the root cause of any flapping event is investigated and rectified before the incident can be declared as resolved. If a user tries to manually set an incident with flapping events to resolved, they are warned that flapping events might cause the incident to reopen.

IBM Cloud App Management provides a built-in event policy called **Global flapping detection** to identify flapping events. The policy detects events that clear and reopen 4 or more times in an hour, and marks them as flapping. If these events stop changing states for more than 30 minutes, they are no longer considered to be flapping. This policy applies to all events and is enabled by default. To view this policy, go to the IBM Cloud App Management **Administration** page, click **Policies**, and ensure you are on the **Event policy** tab. Look for **Global flapping detection** in the list of event policies.

You can change the built-in policy to customize it to your environment. For example, you might want to make the policy more sensitive to recurring problems by decreasing the frequency of the state changes within the same time period before the event is considered as flapping. To do this, you can set the **Number of state changes** to 2 within an hour, meaning if an event changes state twice or more within an hour, it is marked as a flapping event.

You can also set up separate event policies to detect flapping events. For example, you might want to detect flapping events from specified systems such as servers hosting web applications. When the host experiences spikes in CPU loads, the resulting warning events might open and clear, and then repeat again and again. This might point to a hanging process using a large portion of the host's processing power from time to time. You can set up a flapping event policy to detect such events for the hosts to ensure any problem receives the right attention immediately.

To set up this policy:

### Procedure

1. Click **Policies** on the IBM Cloud App Management **Administration** page.
2. Click **Create event policy**.
3. Go to **Details** and enter a name in **Policy name**, for example, Detect flapping from webapp hosts. You can also add an explanation of the policy in **Description** to help you and others understand the purpose of the policy, for example, Detect flapping events from webapp hosts to prevent incidents being resolved until root cause of event is rectified.
4. Click **Specify conditions** in **Events**, and set **Condition 1** as follows: select **Hostname** from the list of attributes, select **contains** from the list of operators, and enter webapp in the field.

**Note:** This example assumes there is a naming convention in place for host names serving web applications, with all such host names having webapp included.

5. Optional: When selecting **Specify conditions**, you can check to see how many events would have matched the conditions you set. Go to the end of the **Events** section, select the number of days between 1 and 30, and click **Test**. The result shows how many events would have matched the policy conditions.

Click **Show results** to view a list of all the events that would have matched the conditions in the set time. Click **New test** to change the time frame for testing, or if you changed conditions and want to check again for matching events.

**Note:** If your event policy enriches fields used by the conditions of your policy, you might not find any matching events after the policy is enabled and applied.

6. Select the **Flapping** check box in **Action**, and expand the section.
  7. Set the fields as follows:
    - a) In the **Enter flapping state** section, go to the **Number of state changes** field and use the arrows to select a value of 4, and then set the **Time period** to 5 Minutes.
    - b) In the **Exit flapping state** section, set the **Time period** to 10 Minutes.
- This means that if an event changes state 4 or more times within 5 minutes, it is considered to be flapping. If the event does not change state for more than 10 minutes, it is no longer considered to be flapping, and will not prevent incident resolution.
8. Set **Enable** to **On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
  9. Click **Save**.


## Managing incident policies

Cloud App Management has built-in incident policies that assign a priority to incidents based on the event severity. You can also define policies to take actions on newly generated incidents to automate some of the incident handling. For example, you can build a policy that assigns certain types of incidents to specific groups or users, and notify others of them automatically.

### Procedure

To create an incident policy, complete the following steps:

1. From the Cloud APM console menu bar, select **Administration** and, in the page that opens, click **Policies**.



2. Click **Create incident policy**.
3. Enter a name and a description for the policy in **Details**.
4. Specify the incidents that you want the policy to apply to in **Incidents**. You can specify to have all incidents considered for the policy actions by clicking **All incidents**, or you can configure what conditions the incidents have to meet before the actions are applied to them by clicking **Specify conditions**.

**Tip:** When selecting **Specify conditions**, you can join multiple conditions using the AND and OR operators. You can also use the example conditions provided by clicking **Use example**. To view the examples, expand **Information and examples** > **Show examples**. In addition, you can select from a list of predefined conditions to use by clicking **Add predefined condition**.

5. Optional: When selecting **Specify conditions**, you can check to see how many incidents would have matched the conditions you set. Go to the end of the **Incidents** section, select the number of days between 1 and 30, and click **Test**. The result shows how many incidents would have matched the policy conditions.  
Click **Show results** to view a list of all the incidents that would have matched the conditions in the set time. Click **New test** to change the time frame for testing, or if you changed conditions and want to check again for matching incidents.
6. Go to the **Action** section and set the actions that you want the policy to take against the incidents.
  - **Assign and notify:** Automatically make a group or user owner of the incident when the set conditions are matched. Select the groups, users, or integrated tools such as Slack to notify about the incidents. If groups are selected to be notified, all members of that group are notified.

- **Set priority:** Set the priority level for incidents that meet the condition, determining how important the incidents are. Incident priority ranges from 1 to 5, with 1 being the highest priority. The priority of the incident is based on the severity of the events that make up the incident, with the highest severity event determining the overall priority of the incident.

**Important:** By default, built-in incident policies set the priority of incidents as described in “[Events and incidents](#)” on page 728. The built-in incident policies are enabled by default. They are called **Set Priority number**, for example, **Set Priority 1**, and range from 1 to 5. To view the built-in policies, go to the **Incident policy** tab on the **Policies** page. You can create new policies that set the priority for incidents or modify the built-in default policies. Understand your requirements before modifying the built-in policies.

7. Set **Enable to On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
8. Click **Save** to save the policy and return to the policy list.
9. If necessary, adjust the order in which the policies are applied with the options in the **Actions menu** for moving a policy up or down the list.  
If a policy has a conflicting rule with one that comes earlier in the list, the rule of the policy that comes after overrides the earlier one.

### What to do next

For an example of creating a policy to assign incident priority, see [Create an incident Policy in the scenario, “Getting Started: Proactively manage the health of your application environment – regardless of size” on page 34](#). See also the example topics that follow for other conditions and actions that you can use to set up policies against incidents.

## Example: Assigning top priority incidents automatically

You can create incident policies to automatically assign specific incidents to a group or a user. You can also send notifications about the incident to groups, users, and tools such as Slack.

### About this task


In this example, you want all high priority incidents that include events from the WebSphere MQ resources to be automatically assigned to the WebSphere MQ administration group. At the same time, you want to notify the group's team leader of such incidents. Setting up this policy helps route incidents to the right personnel efficiently.

Incident priority ranges from 1 to 5, with 1 being the highest priority. The priority of the incident is based on the severity of the events that make up the incident, with the highest severity event determining the overall priority of the incident. For example, if an incident contains critical severity events, then the incident priority is set to 1, the highest priority level. This is the default behavior, and is based on a set of built-in incident policies that set the priority of incidents. Adding new policies or modifying the built-in policies changes the default behavior. For more information, see [“Events and incidents” on page 728](#).

### Procedure

Complete these steps to define a policy for assigning a group to incidents for WebSphere MQ resources:

1. From the Cloud APM console menu bar, select **Administration** and, in the page that opens, click **Policies**.



2. Click **Create incident policy**.
3. Go to **Details** and enter a name in **Policy name**, for example, **Assign high priority MQ incidents to WMQ admins**. You can also add an explanation of the policy in **Description** to help

you and others understand the purpose of the policy, for example, Automatically assign any priority 2 or higher incidents from WebSphere MQ to the WMQ admin group, and notify team leader.

4. Click **Specify conditions** in **Incidents**, and set the following conditions:
  - a) Go to **Conditions > Incident has the following attributes** and set the incident attribute as follows: select **Priority** from the list of attributes, select **is higher than or equal to** from the list of operators, and select **2** from the list of priority levels.
  - b) Ensure you have **AND** set.
  - c) Go to **Describe the events that the incident contains** and click **Add condition to describe incident events**.
  - d) Set **Condition 1** as follows: Select **Resource type** from the list of attributes, select **contains** from the list of operators, and enter mq in the field.
5. Optional: When selecting **Specify conditions**, you can check to see how many incidents would have matched the conditions you set. Go to the end of the **Incidents** section, select the number of days between 1 and 30, and click **Test**. The result shows how many incidents would have matched the policy conditions.  
Click **Show results** to view a list of all the incidents that would have matched the conditions in the set time. Click **New test** to change the time frame for testing, or if you changed conditions and want to check again for matching incidents.
6. Select the **Assign and notify** check box in **Action**, and expand the section.
7. Click **Add assignment / notifications**.
8. On the **Groups** tab, select the WebSphere MQ administration group in the **Assign** column.
9. Go to the **Users** tab and select the check box for the group's team leader in the **Notify** column.
10. Set **Enable** to **On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
11. Click **Save** to save the policy and return to the policy list.

## Results

When priority 2 or higher incidents are created based on events received from WebSphere MQ resources, the incidents are automatically and immediately assigned to the WebSphere MQ administration group to take action. Each group member receives an email with the option to either investigate the incident, or to assign the incident to themselves straight away. In addition, the group's team leader is notified to keep track of such high importance issues.

## Example: Escalating incidents automatically

You can set up incident policies to escalate incidents to selected users, groups, or integrated tools such as Slack channels.

### About this task

For example, you can ensure that the highest priority incidents are investigated without delay by setting up notifications to the right channels and users. Following on from the example described in [“Example: Assigning top priority incidents automatically”](#) on page 743, you can create escalation rules for any high priority incidents from your DB2 servers.

With high priority incidents automatically assigned, each group member receives a notification, with the email providing buttons to take action to investigate or assign the incident to themselves. If the incident is not set to in progress within 15 minutes of the incident being created, you can add an escalation rule to send a notification to the DB2 admin group's Slack channel to highlight that a high priority incident is still awaiting action. If after a further 10 minutes the incident is still not in progress, you can set a rule to notify the team leader that the incident is still open. The team leader can then take action and select a user to investigate.

To set up this policy:

## Procedure

1. Go to **Administration > Policies**, and click the **Incident policy** tab.
2. Look for the incident policy in the list titled "Assign high priority DB2 incidents to DB2 admins". This is the policy created in ["Example: Assigning top priority incidents automatically"](#) on page 743.
3. Click the **⋮ icon > Edit** in the row for the policy.
4. Go to **Action** and expand the **Assign and notify** section.
5. Click **Add escalations** in the **Escalations** section.
6. Click the **Integrations** tab and select the check box for the DB2 admin group's channel in the **Escalate** column. Click **Apply**.
7. Set the **Escalate after:** field to 15 minutes.
8. Click **Add escalations** again.
9. Click the **Users** tab and select the check box for the group's team leader in the **Escalate** column. Click **Apply**.
10. Set the **Escalate after the previous escalation** field to 10 minutes.
11. Set **Enable** to **On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
12. Click **Save**.

## Results

When high priority DB2 database incidents are not set to in progress by any user within 15 minutes of the incident being created, the DB2 admin group's Slack channel receives a notification to remind the group. If after a further 10 minutes the incident is still not in progress, the team leader receives an escalation to ensure action is taken.

**Note:** For the notification and escalation features to work, you must have the email and mobile phone details added for users, as described in ["Defining users and groups"](#) on page 761. In addition, for notifications and escalations to outgoing integrations such as Slack to work, you must have the integration with the third party tool configured as described in ["Configuring outgoing event destinations"](#) on page 717.

## Example: Setting incident priority

If you have a mission-critical data center that provides essential services for your operations, you can change the way incidents are prioritized from that data center. You can create a policy to set higher priority to the incidents from the data center than they would have based on the built-in default settings.

### About this task

Incident priority ranges from 1 to 5, with 1 being the highest priority. The priority of the incident is based on the severity of the events that make up the incident, with the highest severity event determining the overall priority of the incident. By default, the built-in **Set Priority** incident policies rank the incidents in importance as follows:

- Priority 1: if an incident contains critical severity level events.
- Priority 2: if an incident contains major severity level events.
- Priority 3: if an incident contains minor severity level events.
- Priority 4: if an incident contains warning severity level events.
- Priority 5: if an incident contains information or indeterminate severity level events.


You can change how the priority is determined for incidents from the data center by adding a policy that sets any incident that contains major severity events to be a priority 1 incident, ensuring that issues receive attention more quickly even if they do not yet contain critical events.

This example assumes the data center has the **Location** attribute in the events set to NewYork

## Procedure

Complete these steps to define the policy:

1. From the Cloud APM console menu bar, select **Administration** and, in the page that opens, click **Policies**.



2. Click **Create incident policy**.
3. Go to **Details** and enter a name in **Policy name**, for example, Set priority 1 for data center incidents. You can also add an explanation of the policy in **Description** to help you and others understand the purpose of the policy. Example, Set incident priority level to 1 for major events from data center NewYork.
4. Click **Specify conditions in Incidents** and set the following conditions:
  - a) Go to **Conditions > Incident has the following attributes** and set the incident attribute as follows: select **Priority** from the list of attributes, select **is higher than or equal to** from the list of operators, and select **5** from the list of priority levels.
  - b) Ensure you have **AND** set.
  - c) Go to **Describe the events that the incident contains** and click **Add condition to describe incident events**.
  - d) Set **Condition 1** as follows: select **Location** from the list of attributes, select **is** from the list of operators, and enter NewYork in the field.
  - e) Click **Add condition** and ensure you have **AND** set.
  - f) Select **Severity** from the list of attributes, select **is greater than or equal to** from the list of operators, and select **Major** as severity.
5. Optional: When selecting **Specify conditions**, you can check to see how many incidents would have matched the conditions you set. Go to the end of the **Incidents** section, select the number of days between 1 and 30, and click **Test**. The result shows how many incidents would have matched the policy conditions.  
Click **Show results** to view a list of all the incidents that would have matched the conditions in the set time. Click **New test** to change the time frame for testing, or if you changed conditions and want to check again for matching incidents.
6. Select the **Set priority** check box in **Action**, and expand the section.
7. Go to **Set the priority for the incidents described above** and select **Priority 1**.
8. Set **Enable to On** to start using the policy. The policy might take up to 30 seconds to become active and its settings to take effect.
9. Click **Save** to save the policy and return to the policy list.

## Results

When incidents from the NewYork data center arrive containing major severity events, the priority for those incidents is changed to the highest priority instead of setting them to priority 2. This can ensure that problems occurring at the data center are acted upon before they become critical issues, thus helping to avoid disruptions to service.

## Managing incidents

The **Incidents** tab gives you a list of your current incidents. You can view all incidents, or incidents that are assigned to you or groups you are a member of. You can take ownership of incidents, and work with your teams and tools to resolve incidents.

### About this task

Overview of the **Incidents** tab.






Region	Description
1	<p>Incident lists</p> <ul style="list-style-type: none"><li>• <b>My incidents:</b> You can view incidents that are assigned to you.</li><li>• <b>Group incidents:</b> You can view incidents that are assigned to groups you are a member of.</li><li>• <b>All incidents:</b> You can view all incidents.</li></ul> <p>Incidents are sorted based on priority level and the last time they changed, with the highest priority and the latest incident to have changes shown at the top of the list. Incidents of all priority levels are displayed by default.</p>
2	<p>Search and filter fields</p> <p>Use the <b>Search</b> field to find incidents. You can use spaces when searching for more than one word, for example, when searching for a specific incident description.</p> <p>Use the  <b>Filter</b> to display incidents that do not have an owner, or to display incidents in specific states, such as <b>Unassigned</b> or <b>In progress</b>. You can also filter for incidents based on their priority level. Incidents of all priority levels are displayed by default.</p> <p><b>Note:</b> Select <b>No owner</b> to display all incidents that do not have a user assigned as the owner, even if the incident is assigned to a group.</p>

Table 96. Incidents tab overview (continued)

Region	Description
3	<p>Incident summary</p> <p>Displays information about the incident, including ID, priority level, short description, and ownership. Also shows the time the incident was last changed, and how long the incident has been open for based on the time elapsed since the first occurrence of the associated events. The <b>Open for</b> label changes to <b>Duration</b> when the incident is set to resolved.</p> <p>The incident description is based on the resource data contained in the event information. The same resource data is used to correlate the events into an incident.</p> <p>You can take ownership of incidents or assign them to other groups or users by clicking  <b>Menu overflow</b> &gt; <b>Assign</b>. You have the option of assigning the incident to a group you are a member of, or to another user who is a member of that group. You can also click <b>Show all</b> to have all groups displayed, and assign the incident to a group you are not a member of. If the incident is assigned to a group already, but not to a user within that group, then all groups are displayed. Alternatively, click the <b>User</b> tab to look for a specific user to assign the incident to. If you click <b>User</b> and select a user who is a member of more than one group, then you must specify which group the incident is assigned to.</p> <p>You can also resolve an incident here by clicking  <b>Menu overflow</b> &gt; <b>Resolve</b>.</p>
4	<p>Incident bar</p> <p>Displays the icon for the highest event severity level that occurs in the incident, together with a total count for such events.</p> <p>On the left, a link shows the total number of events that are part of the incident. Clicking the link opens the <b>Events</b> tab of the incident details page. On the right, a link opens the <b>Resolution view</b> where you can investigate the incident in more detail, and includes options for resolving it.</p> <p>You can also use the grippy  to drag the incident to the sidebar on the right, and assign it to a group or user.</p>
5	<p>Sidebar</p> <p>Shows users or groups, or the incidents assigned to you. Use the drop-down list to switch between them.</p> <p>Drag an incident to a user or group to assign it to them. You can also drag a user or group from the sidebar to an incident to assign the user or group to the incident.</p> <p>Use the grippy  to drag users, groups, or incidents.</p>

Overview of the incident details UI.

**Procedure**

1. To get a bird's eye view of the incidents affecting your operations, see [“Understanding your incidents at a glance” on page 749](#).
2. For an example of how to start managing your incidents, see [“Starting to work with incidents” on page 753](#).



## Understanding your incidents at a glance

Use the Cloud Event Management dashboards to obtain an overview of your incidents, and see if anything requires your attention.

### Operations overview

The **Operations overview** provides an overview of the incidents affecting your operations, with widgets showing an insight into your incidents at a glance.

Click the **Dashboard** tab in Cloud Event Management to access the **Operations overview**.

Click **Open** to access all of the widgets in the dashboard.

### Resources affected by incidents widget

Use the **Resources affected by incidents** widget to see how your environment is affected by incidents of different priority levels. The widget shows an incident count for the selected priorities that have been created against affected applications, services, servers, clusters, and locations. If any of these categories has a new incident in the past hour with the selected priority level, then a red badge in the top right corner of the tile shows the number of resources affected. Hover over the badge to see more detail.

You can filter what priority incidents the widget displays a summary for using the check boxes next to each priority at the bottom of the widget. **Priority 1** is selected by default. You must have at least one

priority selected at all times. You can also set what resources have a summary tile displayed using the **Menu overflow**. All resources are selected by default. You must have at least one resource selected at all times.

Click the tile of a resource type to open a table below listing the resources affected with information about their names, the time past since the last incident was created for the selected priority levels, and a link to information about the incident in the **Resolution view**. Use the search field to find resources by searching for their names.

For example, you might see **3 Applications affected**. Clicking the tile opens a table listing the application details. Hovering over the incident link shows the incident ID and summary. If more than one incident is affecting the resource, then they are listed in a tooltip. You can investigate the incident you are interested in by clicking the incident ID link, or by clicking the **Open in new tab**. The link opens the **Resolution view** for the incident in the same window or in a new tab.

## Resources affected by incidents (All incidents)

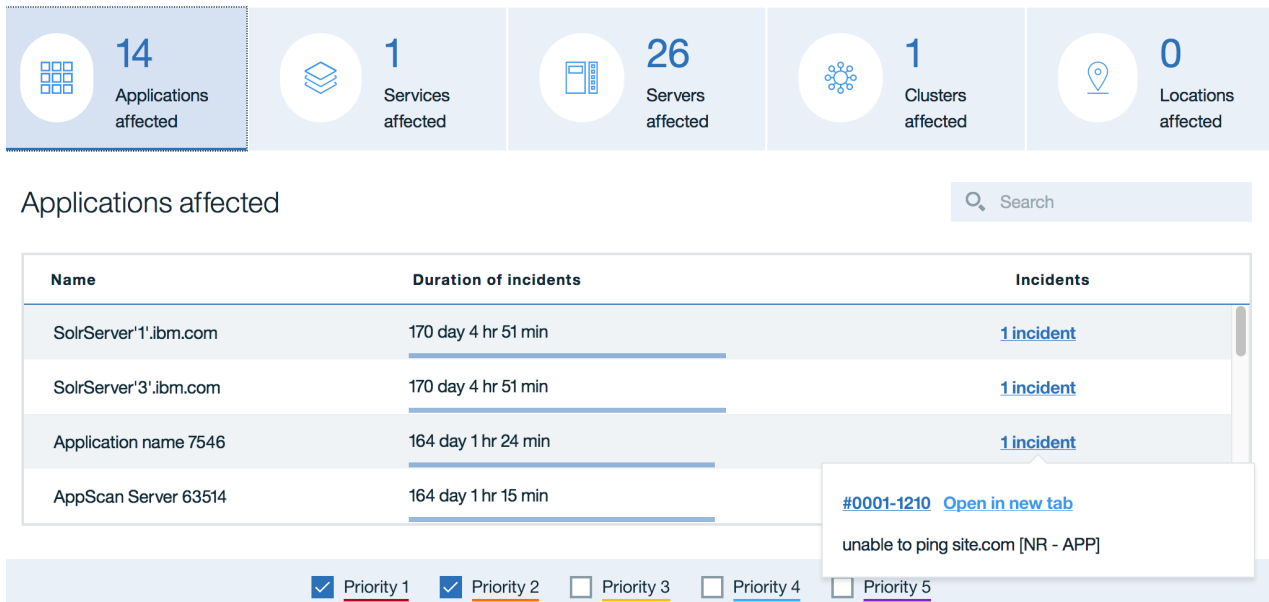


Figure 9. Resources affected by incidents widget

The widget title shows the number of groups selected. Use the **All incidents** drop-down list in the upper right of the window to filter the number of incidents shown according to the groups they are assigned to. If the **Unassigned** filter is selected from the drop-down, and not all groups are selected, then the title also contains **not assigned** to indicate that the count also includes incidents for the selected priority levels that are not assigned to any group. All incidents are selected by default.

### State of incidents widget

Use the **State of incidents** widget to see the high-level status of your incidents, including a total number of incidents for the selected groups, and a count for the different states:

- **Unassigned:** incidents that are not assigned to any group or user. Select the **Unassigned** filter from the **All incidents** drop-down list in the upper right of the window.
- **Assigned to user/group:** incidents that are assigned to a group or to a user within a group, but not set to in progress or on hold.
- **In progress:** incidents that are marked as being worked on.
- **On hold:** incidents for which work has been temporarily suspended.

Click a priority in the legend to have the priority color highlighted for each state and the count for each priority in the various states displayed above the pie chart. If you click the selected priority again in the legend, the highlighting and the count is removed.

You can also hover over a priority color in the pie chart to have the priority color highlighted, and the count for that priority displayed for a single state.

**Note:** Hovering over only works if no priority is selected in the legend.

Click the priority color in the pie chart under any of the incident states to drill down into the details of only those incidents. The **Incident** tab opens displaying only incidents that have the selected priority level and state. You can access further information about each incident from the list.

Use the **All incidents** drop-down list in the upper right of the window to filter the displayed incident counts according to the groups they are assigned to, or show incidents not assigned to any group or user. All incidents are selected by default.

The widget title shows the number of groups selected and their total incident count. If the **Unassigned** filter is selected from the drop-down, and not all groups are selected, then the title also contains **not assigned** to indicate that the count also includes incidents not assigned to any group or user.

### State of incidents (All incidents; 249 incidents)

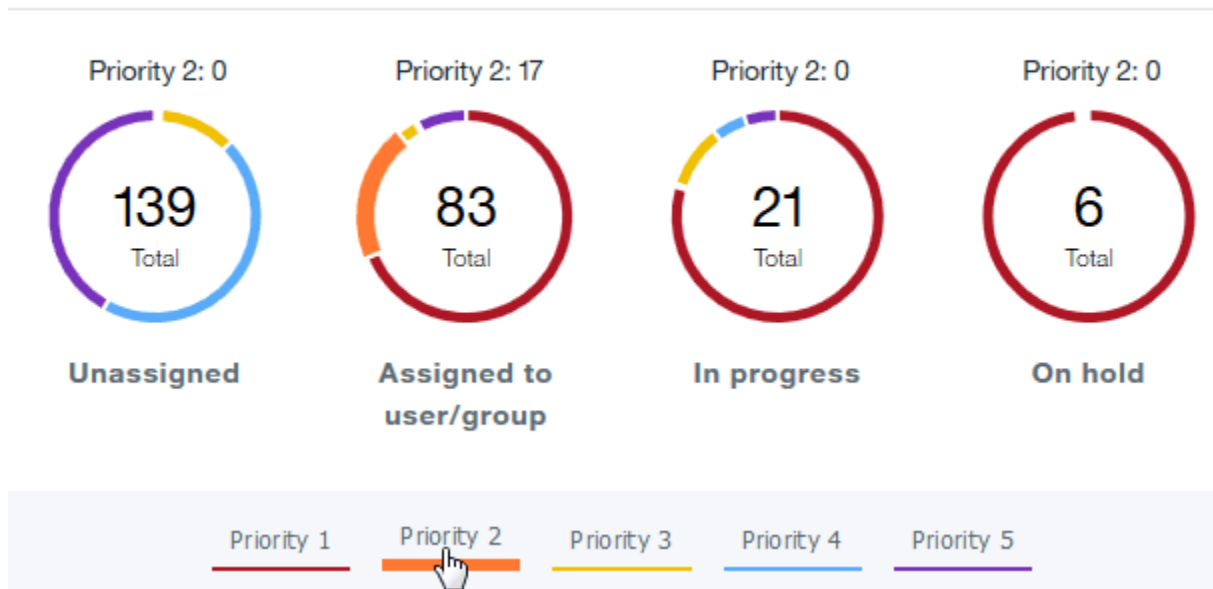


Figure 10. **State of incidents** widget

### Open incidents over time widget

Use the **Open incidents over time** widget to understand the number of incidents created over time. You can select to view incident trends for the last 8 or 24 hours, or for the last 7 days. The number displayed at the end of the trend is the live count for the current number of open incidents based on the latest data. The count is updated every time there is a change in the number of incidents (for example, new incidents are created, or existing incidents are resolved).

Hover over the trend line to see the total number of open incidents for a specific time. The trend line shows the total number of open incidents for every 5 minutes if 8 or 24 hours is selected, or for every hour if 7 days is selected.

**Note:** If data is not available for the full 8 or 24 hours, or for all of the last 7 days, the trend line is not displayed for the full time period. For example, if incident data is only available for the last 5 days, and 7 days is selected to be displayed, then the trend line only displays the 5 days worth of data available.

You can also set a filter to only show trends for specific incident priorities, or see trends for all priorities. All priorities is selected by default.

The widget title shows the number of groups selected. Use the **All incidents** drop-down list in the upper right of the window to filter the number of incidents according to the groups they are assigned to. If the **Unassigned** filter is selected from the drop-down, and not all groups are selected, then the title also contains **not assigned** to indicate that the trend data also includes incidents not assigned to any group. All incidents are selected by default.

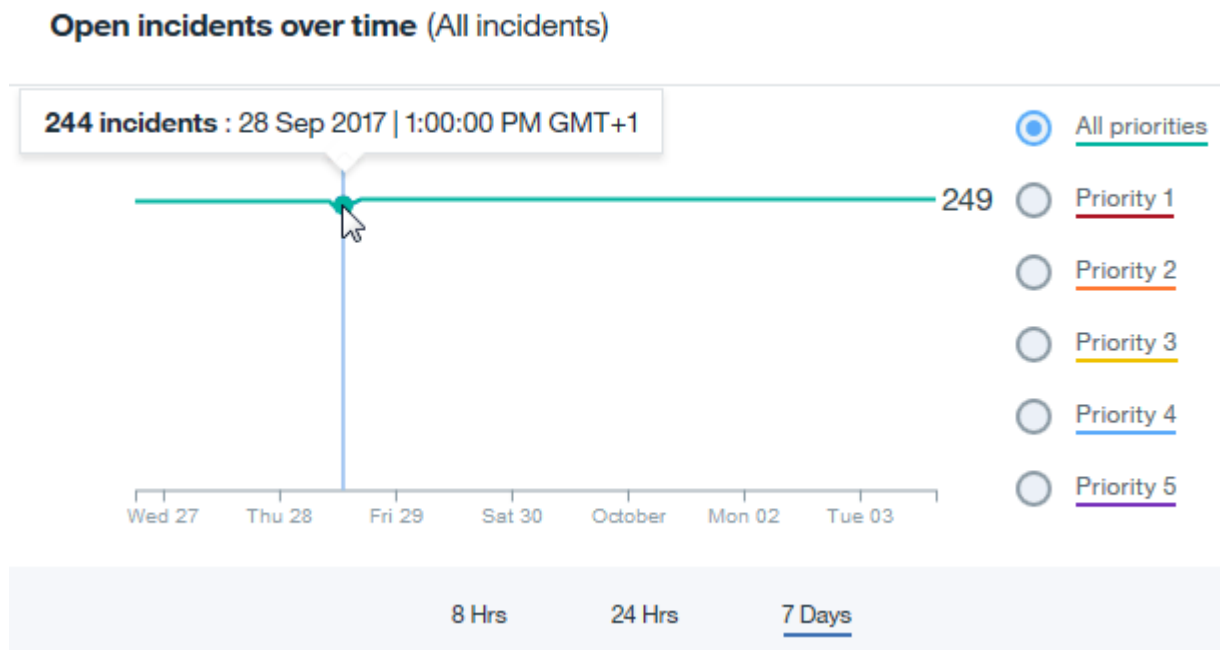


Figure 11. *Open incidents over time* widget

## Understanding the mean time to resolve incidents

### Efficiency Overview

The **Mean time to resolve and respond** dashboard provides an overview of the mean time to resolve incidents within your operations, with widgets providing insight into your incidents at a glance.

Click the **Dashboard** tab in Cloud Event Management to access the **Mean time to respond and resolve** dashboard.

Click **Open** to access all of the widgets in the dashboard.

### Average duration widget

Use the **Average duration** widget to see the mean time to resolve your incidents, the mean time to respond to an incident, and the number of closed incidents. You can filter the incidents by priority level starting from a **Priority 1** to a **Priority 5**. You can view average durations by 24 hours, 7 days, 30 days, and 90 days. By default the graph will show the average duration for the last 30 days. There is no default priority setting.

There are three tiles displaying important information about the operational efficiency of your incidents. The **Mean time to incident resolution** metric displays the mean time from incident generation to resolution. The **Mean time to respond to an incident** metric displays the mean time from incident generation to in progress for the first time. The **Number of closed incidents** metric is a count of the number of incidents closed during a certain time period. The **Time an incident is on hold** metric is a running total of the time all incidents within a certain time period are left on hold. The **Opened incidents** metric is a count of opened incidents created during a certain time period.



Figure 12. **Average duration** widget

The **Average duration** widget will only display the minimum and maximum values in the **Opened incidents** column.

You can change the data displayed in your chart depending on the options that you select. By default the mean time to resolution, mean time to respond, unresolved incidents, time an incident is on hold are automatically checked and are set display data for the last 30 days. Click the **Filter by** and **Priority** icon to sort and view incident priority.



## Starting to work with incidents

Learn how to start managing incidents with IBM Cloud App Management.

### About this task

If your IBM Cloud App Management set up is ready, you can start managing incidents. The following is an example of how to access your incidents and start investigating them.

### Procedure

1. Go to the **Incidents** tab of the em\_start.html user interface.
2. View incidents that are assigned to you on the **My incidents** tab. Administrators and other users might have already assigned incidents to you. If you do not have any incidents that are assigned to you, click **Go to group incidents** to see what incidents are assigned to groups you are member of.
3. Click  **Filter** and select **No owner** to show incidents that do not have a user assigned as an owner yet. This filter also shows incidents that have been assigned to a group, but not to a user.
4. Take ownership of incidents by dragging the incident to the sidebar on the right, or by clicking  **Menu overflow** > **Assign**.
5. Click **My incidents** and click **Events** to learn more about the events that make up the incident. The **Events** tab opens.
6. On the **Events** tab, investigate the information available about the most severe events related to the incident. The top level information in the table shows data such as event severity, the type of resource

sending the event (for example, application or server), when the event first occurred, a summary describing the event, and the type of event in short. By default, events are sorted based on severity, with the highest severity at the top of the list. You can change the sort order by clicking the column headers.

Expand the row for an event to find out more about the problem, such as what state the event is in, the host affected, the URLs for the monitoring system sending the event data, the number of times such an event has occurred (count), and other details. You can click the **See more info** button to access all details available for the selected event.

7. Click the **Timeline** tab to see the history of the incident. You can see that another user from your group posted a comment. The comment suggests a similar problem occurred not long ago, and the user has notes about what steps were taken to resolve the problem.
8. Go to the **Resolution view**, find the user in the **Collaborate** column, and click **Notify**. Enter a message and ask the user to share their notes, then click **Send**.
9. Set the incident status to **In progress**.

## Resolving incidents with runbooks


Runbooks provide structured steps to help solve incidents.

### Before you begin

To have runbooks available to use for your events, you must first define runbooks as described in “[Managing runbooks](#)” on page 764, and then set up event policies where runbooks are associated with events as described in “[Setting up event policies](#)” on page 727.

The following is an example of how to use runbooks to address the events that form an incident, and as a result resolve the incident itself.


### Procedure

1. Go to the **Incidents** tab of the Cloud Event Management user interface.
2. Go to **My incidents** and click **Investigate** to retrieve more information about the incident. The **Resolution view** displays suggested runbooks for the type of incident.
3. In the **Resolution view**, click  **Menu overflow** > **Run** next to the runbook you want to apply.

If the runbook uses parameters, the parameter values are based on the event policy, and depend on the events associated with the selected runbook:

- If there is only one event, or if there are multiple events all with the same parameter values, then the parameter values for the runbook are taken from a single event, and the runbook is launched using those values.
- If multiple events with different parameter values are correlated into an incident, each event's parameter values are displayed. Select the value you want to run the runbook against and click **Run**.

The Runbook Automation UI is displayed where you can work with the runbook. For more information, see [https://www.ibm.com/support/knowledgecenter/SSZQDR/com.ibm.rba.doc/CR\\_runrunbook.html](https://www.ibm.com/support/knowledgecenter/SSZQDR/com.ibm.rba.doc/CR_runrunbook.html).

**Tip:** You can also apply runbooks associated with the events from the **Events** tab. Click the **Events** tab, expand the row for the event, and click **Suggested runbooks** to view the available runbooks. You can click  **Run** for the runbook you want to apply. Parameters values for the runbook are derived from the event, or you might be prompted to enter a value manually either as it requires information such as a user name, or the runbook is set up to request the value at runtime.

For more information about viewing the available runbooks, reviewing the runbooks that you have used to date, and running the runbooks, see [https://www.ibm.com/support/knowledgecenter/SSZQDR/com.ibm.rba.doc/WR\\_workRunbooks.html](https://www.ibm.com/support/knowledgecenter/SSZQDR/com.ibm.rba.doc/WR_workRunbooks.html).

4. The runbook completes and solves the underlying problem causing the incident. The events that formed the incident are then cleared, and in turn the incident is automatically set to resolved and closed.

### What to do next

For information about creating and managing runbooks, see [“Managing runbooks” on page 764](#).

## Managing thresholds


Thresholds test for resource issues such as a slow response time. When the conditions of a threshold are true, an event is opened and an incident is generated. You can create, edit, delete, enable, or disable thresholds.

### Procedure

Open the **Threshold Management** page to view and manage thresholds:

1. Click **Administration > Thresholds**.

A table of defined thresholds is displayed:






- **Name** is the title given to the threshold when it was saved. Click the name to view and change the definition in the threshold editor page.
  - **Severity** is the severity that was chosen for the threshold.
  - **Assigned to** is the resource type that the threshold is defined to monitor, such as Linux Systems or Kubernetes Service.
  - **Permissions** are either "Read-only" or "Editable". Read-only thresholds are predefined or imported from integrated sources and cannot be changed. Editable thresholds were created by a member of your team and you have full editing capability.
  - **State** is "Enabled" for thresholds that are operational, which means they are monitoring the resources that they were assigned to. **State** is "Disabled" when **Enable** has been turned off and the threshold is non-operational. (See steps [“3” on page 755](#) and [“9” on page 757](#).)
2. If you're looking for a threshold that doesn't show on the first page of the table, use the page controls:
    - Click inside the Filter text box and type the beginning of the value to filter by. As you type, the rows that do not fit the criteria are filtered out. For example, type begin typing Indeterminate to filter the list down to only those thresholds with severity Indeterminate.
    - Select the column to sort by.
    - Select the number of thresholds to show per page: 10, 25, or 50.
    - Select the next or previous page or a specific page number.
  3. Complete one of the following steps:
    - To define a new threshold, click **Create**. Continue to step [“4” on page 755](#).
    - To edit a threshold definition, click the threshold name. Continue to step [“4” on page 755](#).
    - To delete a threshold (or more), select its check box and click **Delete** in the banner that appears. After you respond to the confirmation prompt by clicking **OK**, the threshold is permanently deleted.
    - To enable or disable a threshold, click the  icon and select **Enable** or **Disable** from the list.

Create or edit the threshold definition:

4. Complete the **Details** section:

- a) **Threshold name** must start with a letter and can have up to 63 letters, numbers, and underscores.
- b) Enter a description for the threshold.  
As well as in the **Thresholds** pages, the description is displayed in the **Incident** details (select the **n Events** link followed by **See more info**).

5. Complete the **Threshold** section:

- a) For **Type of resource to create the threshold on**, select the resource type that you want to monitor, such as Linux Systems.
- b) For **Threshold severity**, select  critical,  major,  minor,  warning, or  indeterminate.
- c) For **Consecutive samples**, specify how many consecutive threshold samples must evaluate to true before an event is generated: A threshold with a setting of 1 and a sample that evaluates to true, an event is generated immediately; a setting of 2 means that two consecutive threshold samples must evaluate to true before an event is opened.

d) Define the condition:

- 1) Select the metric to compare from the metric list. (The remaining fields vary depending on the type of metric.)
- 2) If the relational operator field is displayed, select one: < less than, <= less than or equal, = equal, >= greater than or equal, > greater than, or != not equal.
- 3) If this is a text metric, you can also select one of these relational operators:




**MISSING** to enter a list of text entries to compare. If none of the entries matches the data sample when the threshold is evaluated, an event is opened.

**match** or **not match** to enter a regular expression to compare. The **match** and **not match** operators look for a pattern match to the expression. If the regular expression matches or does not match the data sample when the threshold is evaluated, an event is opened. The easier it is to match a string with the expression, the more efficient the workload at the managed resource. The expression does not need to match the entire line; only the substring in the expression. For example, in "See him run", you want to know if the string contains "him". You could compose the regular expression using `him` but you could also use `.*him.*`. Or, if you are looking for "See", you could enter `See`, or you could enter `^See` to confirm that it's at the beginning of the line. Entering `.*` wildcards is a less efficient search and raises the workload. For more information about regular expressions, search for "regex" in your browser.

- 4) Enter the value to compare using the allowed format for the metric, such as 20 for 20% or 120 for 2 minutes.

For example, a threshold condition of `Process Percent User Time > 5%` tests if the metric sample for Process Percent User time is greater than 5% and opens an event if the comparison is true.

e) Optional: Add another condition to the threshold:

- 1) Select **Add condition** or **Add nested condition** (see [Example](#)).
- 2) Leave the logical operator at  AND if the previous condition and this condition must be met for the threshold to be breached or, if either of them can be met for the threshold to be breached, toggle the  AND button to  OR.
- 3) Select the metric to compare from the **Metric condition** list.
- 4) Select the relational operator: < less than, <= less than or equal, = equal, >= greater than or equal, > greater than, or != not equal.
- 5) Enter the value to compare using the allowed format for the metric.

If you are adding multiple conditions to a threshold or adding a display item (step "6" on page 757), select metrics from the same metric list (data set). Otherwise, you might get an error message while defining the threshold.

- f) Optional: Add an aggregation expression that applies to the data that meets the defined condition (or conditions):
  - 1) Select the aggregation metric from the list.
  - 2) Select **average** for numeric metrics, **count** for text metrics, or **none**.



- 3) Select the relational operator: < less than, <= less than or equal, = equal, >= greater than or equal, > greater than, or != not equal.
  - 4) Enter the value of the aggregation metric.
6. Optional: Select a **Display item** if one is available and you want to continue evaluating the threshold on other data sample rows.

After a row evaluation causes an event to open, no more events can be opened for this threshold on the monitored resource until the event is closed. By selecting a display item, you enable the threshold to continue evaluating the other rows in the data sampling and open more events if other rows qualify. Display item is not available if the threshold includes an Aggregation condition.

**Known limitation:** If you deploy the runtime data collectors in on-premises environment, when you define the threshold and select **Display item**, metrics of the selected item might not display. Best practice is to not select a display item for data collectors in an on-premises environment.

7. Select the resources for the threshold to monitor in the **Assignments** section:
- Select **All resource\_type** to apply the threshold to all resource instances of the same type, such as all Hadoop hosts.
  - Select **Individual instances** to see and select the resource instances. Individual instances cannot be selected for the WebSphere Applications agent nor any other agent that has subnodes.
  - Select **Group(s)** to see and select from the list of resource groups. (For more information, see [“Managing resource groups” on page 758.](#)) If you assign the threshold to a resource group that does not include a managed resource of the threshold's resource type, a message notifies you that the threshold does not apply to any resource types in the group.
8. Available only for thresholds that are created for and assigned to Linux OS systems: Complete the **Define reflex action** section if you want to execute a command when an event is opened.

a) Enter the command to execute.

In this example, two commands are run by the Linux OS agent: The text in quotes is echoed and redirected to a log file, and the `clean_logs` script runs on the associated Linux OS disk (&{KLZ\_Disk.Disk\_Name} is replaced by the attribute value).

```
echo "`date` : WT_LZ_user_login is true for &{KLZ_Disk.Disk_Name}"
>>/tmp/wt.log;/scripts/clean_logs.sh &{KLZ_Disk.Disk_Name}
```

When you want to include an attribute value in the command and don't know the exact spelling, you can look up the dataset name and attribute name in the Attributes chapter of the [Linux OS agent Reference](#) [\[2\]](#).

- b) Select one of these options to control how often the command is run:
- Select **On first event only** if the data sample has multiple rows and you want to run the command for only the first event occurrence in the data sample. Clear the check box to run the command for every row that causes an event.
  - Select **For every consecutive true interval** to run the command every time the threshold evaluates to true. Clear the check box to run the command when the threshold is true, but not again until the threshold evaluates to false, followed by another true evaluation in a subsequent interval.
9. If you don't want the threshold to begin monitoring, drag the **Enable** slider from On to Off. The thresholds table shows "Disabled" in the **State** column.

## Results

After you save the threshold, it starts on the resources instances that were assigned in step [“7” on page 757.](#)

## Example

Nested conditions are used to support multiple conditions joined with mixed AND and OR operators. Otherwise, multiple conditions would use Boolean AND logic or Boolean OR logic, not both. To illustrate,

the following threshold evaluates to true if either the process CPU is greater than ½ second and the process command is named kynagent or if the process command is named klzagent:

Condition 1	Process CPU Seconds >= 0.5 seconds AND Process Command Name = kynagent
Condition 2	OR Process Command Name = klzagent

The intention, however, is for the threshold to evaluate to true if the process CPU is greater than ½ second and the process command is named either kynagent or klzagent. To achieve the desired result, select **Add nested condition** for Condition 2:

Condition 1	Process CPU Seconds >= 0.5 seconds AND
Condition 2 (nested)	Process Command Name = kynagent OR Process Command Name = klzagent

### What to do next

- View, edit, disable or enable, or delete the threshold in the **Threshold Management** table
- Follow this usage scenario to get some hands-on practice with creating thresholds. See [“Getting started: Accelerate your transition to the cloud with DevOps”](#) on page 38
- Start monitoring your resource as described in [Monitoring resources in your environment](#)

## Managing resource groups

Your monitored environment might have multiple managed resources that can be categorized by their purpose. Such resources often have the same threshold requirements. Use the **Resource groups management** page to organize managed resources into groups that you can assign thresholds to. You can assign thresholds to resource groups for monitoring the managed resources of the same type that belong to the group.

### Procedure

Complete these steps to configure and manage resource groups:

1. In the Cloud APM console, select **Administration** and, in the page that opens, click **Resource groups**.



The **Resource groups management** page opens with a list of the configured resource groups. The tags are the resource types of the managed resources in the group.

2. Take one of the following actions: **Create incident policy**.
  - To configure a new resource group, click **Create group**. The **Create resource group** page opens.
  - To edit a resource group, click the resource group name link. Alternatively, you can click **...** and select **Edit**. The Edit resource group page opens.
  - To make a copy of a resource group, click **...** and select **Duplicate**. A copy of the resource group is created, which you can now edit.
  - To delete a resource group, click **...** and select **Delete**. Confirm that you want to permanently delete the resource group when prompted.

3. Configure the resource group:
  - a) Enter a name that starts with a letter and has up to 63 letters, numbers, and underscores.
  - b) Optional: Enter a description of the resource group. The description is useful, especially for other users, to understand the context of the group.
  - c) In the Filters list, select a resource type to see the managed resources of that type.
  - d) In the list that displays, click the plus (+) next to each manage resource that you want to add to the group or click the + for **Add all** resources of that type to the group.  
As you select resources, a pop-up window shows the managed resource name and resource type, and a counter keeps track of how many resources that are in the group.
  - e) Continue to select resource types and add managed resources to add to the group.
  - f) If you want to remove a resource from the group, click the **Delete** button next to it.
4. When you are finished configuring the group, click **Create** or **Save**.

### Results

The **Resource groups management** is displayed with your newly created or edited resource group.

### What to do next

Assign a resource group to a threshold. For more information, see [“Managing thresholds” on page 755](#).

## Setting up users and groups

You must ensure that valid user ids are added to IBM Cloud Private before you can add the users to Cloud App Management.

Then, you can add users to groups, assign user roles, and set up event notifications for the users.

### Logging in to the Cloud App Management UI to authenticate as a user


The *first user* account used to login to IBM Cloud App Management will be assigned the Operations Lead role. This role has administrator privileges. This first user can then create further users.

#### Before you begin

Before you can set up your user profile or create any new users or groups in IBM Cloud App Management, you must have an IBM Cloud Private user account. When you install IBM Cloud Private, a default user account is set up.

Log in to IBM Cloud App Management using a valid IBM Cloud Private user account. The *first user* account used to login to IBM Cloud App Management will be assigned the Operations Lead role. This will provide this user account administrative privilege for the product.

This user account then has the administrative privileges required to:

- Create new users in the IBM Cloud App Management UI.
- If LDAP is used, connect the LDAP directory with the IBM Cloud Private cluster. Import users and groups from the LDAP directory to add to the cluster. For more information, see the [Configuring LDAP connection](#)  topic in the IBM Cloud Private Knowledge Center.

New users created by the *first user* with the Operations Lead (admin) role are sent two emails:

- The first email welcomes the user to the IBM Cloud App Management product.
- The second email requires the new user to validate their email address to the system. Until an email address is validated the account will receive no incident notifications.

#### Note:

This *first user* with the Operations Lead (admin) role does not receive email notifications. Email are only sent to users created subsequently by this user.

## About this task

New users should complete the following steps:

### Procedure

1. Check your email for a welcome email that describes how to log in and verify your email address. This welcome email provides you with your new user ID and it contains links to the **Getting Started**, **Administration**, and **Incidents** pages. The welcome email also contains information about another email, which contains an activation link to enable email notifications. For information about activating notifications via SMS and voice, see [“Activating notifications via SMS and voice” on page 760](#).
2. After you verify your email address, open the second email. This email provides you with a link to access IBM Cloud App Management. You are provided with a link to the **Getting Started** page.

### What to do next

Create users, add users to groups, and assign users to incident policies. For more information, see [Manage users and groups](#).

## Activating notifications via SMS and voice

For IBM Cloud App Management in an IBM Cloud Private environment, notifications via SMS and voice are supported through the Nexmo API.

### Before you begin

A Nexmo account is required.

### Procedure

- Retrieve your API key and API secret from <https://dashboard.nexmo.com>.
- Enable and configure Nexmo in the values .yaml file. The following section must be populated:

```
nexmo:
  # true to use Nexmo, false disables
  enabled: false
  # API key name, from https://dashboard.nexmo.com
  key: ''
  # API key secret, from https://dashboard.nexmo.com
  secret: ''
  # Default Nexmo number from which to send SMS messages
  sms: ''
  # Default Nexmo number from which to send voice messages
  voice: ''
  # Override numbers used for selected countries
  # Property names are country codes, values are objects with "voice" and "sms" properties
  # Enter as a JSON object in quotes
  numbers: '{}'
```

- Mobile phone numbers must be verified before users can receive SMS or voice notifications. A user must complete the following steps to verify a mobile phone:
  - a) In the welcome email notifying a user they have been added to IBM Cloud App Management, click **Edit profile**.
  - b) The user details page is displayed. In the message prompting the user to verify their mobile phone, click **Send verification code**.  
A verification code will be sent to the number provided.
  - c) Enter the code received in the **Verification code** field.
  - d) Click **Save**.

## Defining users and groups

Define users and create groups of users to handle incidents as required.

### About this task

To add users, complete the following steps:

### Procedure




1. Click **Users and Groups** on the IBM Cloud App Management **Administration** user interface.

The list of available users is presented here. If you have users defined, you can expand the section for each user and view information such as:

The groups the user is a member of.

The number of users those groups have.

The number of incident policies associated with each group (for example, policies that automatically assign incidents). Click the user name or the group name to edit their settings.

You can also edit the user details by clicking . Send them a notification by clicking . You can remove users by clicking  in the appropriate row.

2. On the **Users** tab, click **New user**, and enter the following information in the **Details** section:

Option	Description
<b>Full name</b>	Enter the user's full name.
<b>User Id</b>	Enter the User or IBM id that was created in IBM Cloud Private. For more information, see <a href="#">“Logging in to the Cloud App Management UI to authenticate as a user”</a> on page 759.
<b>Email</b>	Enter the user's main email address.
<b>Role</b>	Set one role for each user. The following roles can be assigned: <ul style="list-style-type: none"><li>• <b>Operator:</b> Users have access to incident lists, the dashboard, and their own user profile.</li><li>• <b>Operations engineer:</b> Users have access to incident lists, the dashboard, and their own user profile.</li><li>• <b>Operations lead:</b> Users have full access to all the features of Cloud App Management</li></ul> For more information about Cloud App Management roles, see <a href="#">“Roles”</a> on page 763.
<b>Secondary email</b>	Enter the user's back-up email address.
<b>Mobile phone</b>	To send SMS and Voice message notifications to the user, enter the user's mobile phone number. The mobile phone number format is +[country code][phone number] without spaces or separators. For example: +19585550123
<b>Voice language</b>	Select the language for voice message notifications. The default is US English.

3. Optional: To add the user to one or more groups, expand the **Group** section and click **Assign to group**. Select the check boxes for the groups you want the user to be a member of and click **Assign**.

**Note:** Users can be organized into groups to reflect the structure of your organization and send notifications to multiple contacts at once. For example, you might have UNIX support groups, database administrators, payroll application experts, and customer support teams.

If you do not have the groups that are required, you can create them later and assign users to the groups as described in [“Setting up groups”](#) on page 762.

4. Optional: Expand the **Work hours** section to define the working hours for the user. The settings here are also used in the automatic assignment to shifts when new schedules are created.
  - a. Select a time zone from the drop-down menu. If you have teams in multiple time zones, consider the impact of Daylight Saving Time on each time zone, and the resulting changes in work hours.
  - b. Select the user's working days of the week, and working hours per day.

**Tip:** You can set the working hours for one day and then copy the settings over to other days by clicking the **Duplicate** link after the day you want to copy. Ensure you select the check box for the day you want to copy, and the check box for all the other days you want to copy the schedule to. You can also click **Apply default working hours** to set working hours for the user from Monday to Friday from 8:00 to 17:00.
5. Optional: Expand the **Notify me** section to set the notification preferences for the user. Different notification methods can be selected for normal working hours and off days. You can select more than one option for each:
  - During working hours.
  - When I am not working.
6. Click **Save**.

## Setting up groups

Create groups and add users to groups to organize your teams as required.




### About this task

To create groups, complete the following steps:

### Procedure

1. Click **Users and Groups** on the Cloud App Management **Administration** user interface.
2. Click the **Groups** tab.

The list of available groups is presented here. If you have groups defined, you can expand the section for each group and view information such as the users that are a member of that group, the role of each user, and the incident policies associated with the group (for example, policies that automatically assign incidents). Click the user name or the group name to edit their settings. You can also click the incident policy name to view and edit the policy settings.

You can also edit the group details by clicking . Send the group a notification by clicking . You can remove groups by clicking  in the appropriate row.

3. Click **New group**.
4. On the **Details** tab, set up the group:
  - a) Enter a name for the group in **Group name**.
  - b) Add users to the group by selecting each user from the **User membership** list.
  - c) Select one or more owners for the group from the **Owner** list. Owners are responsible for the administration of the group, but not for resolving incidents. Owners are notified of changes made to the group. Only users that have the **Operations lead** role can make changes to groups.
5. Click **Save**.

## Roles

Cloud App Management provides roles to control the features that users can access.

### Operations Lead

This role grants access to all capabilities of Cloud App Management. Users with this role configure Cloud App Management for their teams, including event sources, policies, integrations with third party tools, and users and groups. They also have full access to incident management capabilities.

The **Getting started** page for users with this role provides general information about Cloud App Management, and links to all capabilities.

**Important:** You must have at least one user that has the Operations lead role assigned. This role provides access to all capabilities, and changing to another role will limit access to Cloud App Management features, including policies and user management.

You cannot change your role if you are the only user with the Operations lead role. A user with the Operations lead role can only change their role if at least one other user has the Operations lead role assigned.

### Operations Engineer

This role grants full access to incident management capabilities, but limits access to parts of the user's profile settings.

This role places the following restrictions on user and group configuration:

- Can view the list of users and groups, but can only edit their own user profile.
- Cannot change their group or role settings in their profile, but can change work hours and notification settings.
- Cannot delete any user or group profile.
- Can send notifications to users and groups.
- Can send notifications to the groups and their members on shift using the **Who is working now?** page. However, they cannot update the shift assignments.
- Cannot edit any group settings.

The **Getting started** page for users with this role provides general information about Cloud App Management, and links to all capabilities.

### Operator

This role grants full access to incident management capabilities, but limits access to all configuration capabilities except to parts of the user's profile settings. In addition, they cannot generate sample events.

Users with this role manage incidents, including resolving them.

They can view the list of users and groups, but can only edit their own user profile. The same user and group configuration restrictions apply as for the **Operations engineer** role.

The **Getting started** page for users with this role provides general information about Cloud App Management, and links to the user's profile, incident lists, and dashboards.

## Using the REST API

IBM Cloud App Management provides a REST API for operations such as sending test events to your IBM Cloud App Management service, managing users and groups for your service, querying incident details, and managing event policies.


### About this task

You need an API key to access your Cloud App Management service. You can then use the key for the API functions.

## Procedure

1. You can use the name and password in your IBM Cloud App Management service credentials to connect to the API, or you can generate a key as follows.
  - a) Click **API Keys** on the Cloud App Management **Administration** tab.
  - b) Click **New API key**.
  - c) Enter a description for the key in **API key description**.
  - d) Specify which part of the Cloud App Management API the key provides access to. Go to **Permissions** and ensure only those check boxes are selected that you want the key to provide access to. All APIs are selected by default.

**Important:** Make a note of the APIs the key provides access to. For example, note it in the description you enter for the key. You cannot view or change which APIs were selected for the key later.
  - e) Click **Generate**. A new name and password are generated. Make a note of these values.

**Important:** The password is hidden by default. To view and be able to copy the password, set **Display password** to **Show**. Ensure you make a note of the password. For security reasons the password cannot be retrieved later. If you lose the password, you must delete the API key and generate a new one.
  - f) Click **Close**.
  - g) Click the  icon next to **Manage API keys** to view the base URL for the API.
2. Use the API key for the API calls to your Cloud App Management service.

You can use the API to create events with custom payloads, trigger sample events, configure users and groups, and query incident properties (for example, retrieve the details of the events that are correlated into a specific incident). For more information about the Cloud App Management API, see [Event Management API documentation](#) in the IBM Cloud API docs.

**Note:** Sample events and their incidents can also be generated from the user interface and viewed in the incident lists as described in [“Managing incident policies”](#) on page 742.

## Managing runbooks

You can create your own custom runbooks and manage your existing catalog of runbooks in IBM Cloud App Management.

### Procedure

1. Click **Runbooks** on the IBM Cloud App Management **Administration** page.

**Important:** You have different levels of access to runbooks depending on your role in IBM Cloud App Management. The **Operations lead** role provides full access to runbook management, including the permission to approve runbooks for publishing. The **Operations engineer** role provides access to runbook management without the publishing approval permission. Users who have the **Operator** role can only preview and run runbooks assigned to them. For more information, see [“Roles”](#) on page 763.
2. Click **New runbook** to create a new runbook.
3. For more information about managing runbooks, including creating runbooks and using sample runbooks, see [IBM Runbook Automation knowledge center](#).



## About data retention and summarization

Learn about the timeline metrics shown in the Cloud App Management console Resource dashboards and how you can configure your IBM Cloud App Management installation to retain historical data for a shorter or longer period of time than the default 8 days and to enable hourly and daily summarization.

### Data sampling

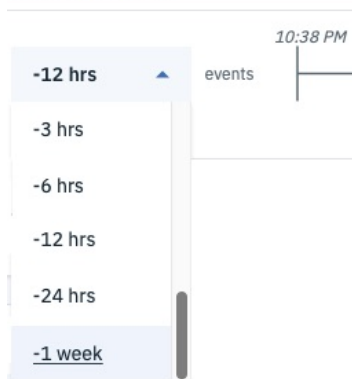
The ICAM Agents and ICAM Data Collectors are monitoring your environment for early detection of performance and availability issues.

Performance data is collected as frequently as once per minute and stored on the Cloud App Management server for 8 days by default. After 8 days, as new data samples arrive, the oldest are removed.

### Resource dashboards

Along with performance and other relevant metrics displayed in the Resource dashboards, you have an Events timeline. The initial display of the timeline and chart views present the past 12 hours. You can adjust the time span to show from 3 hours up to a week. If your Cloud App Management installation is configured for a data retention value of 2 days up to 32 days, the time span options reflect the value that was set.

#### Events timeline



### Data retention and summarization

Data retention is the number of days that data samples are saved before the oldest data samples are deleted to make room for new data samples. The default is 8 days. During Cloud App Management server install or upgrade, you can reconfigure retention to be from 2 to 32 days. Any value beyond 32 days is not recommended and can degrade Cloud App Management performance. If you prefer to keep the prior release's data retention during server upgrade and make a later decision about reducing raw data retention, set **rawMaxDays** to 32. For instructions, see step [“13”](#) on page 147 of *Installing your Cloud App Management server on Red Hat OpenShift*.

Summarization refers to the aggregation of the retained data into time-based categories: hourly for short-term recall; daily for long-term recall. By default, summarization is turned off. During Cloud App Management server installation or upgrade, you can enable summarization (step [“13”](#) on page 147 of *Installing your Cloud App Management server on Red Hat OpenShift*). When summarization is enabled, samples of agent metrics are summarized once per hour for short-term history and once per day for long-term history. The server keeps the summarized metrics for recall: 60 days of short-term history and 6 months of long-term history. You can change the length of time short-term and long-term history data samples are saved by creating a Kubernetes ConfigMap, as described in [“Configuring summarization”](#) on page 766.

Raw data retention offers the greatest flexibility in data visualization and on-demand aggregation. The more days that are retained, the more storage that is needed to store large quantities of raw data. In addition, metric query requests can take longer to fulfill as the number of days of raw data increases.

The retention period for hourly and daily summarized data should be chosen based on your long term trend analysis needs. Although hourly summarized data offers a more granular look at metrics than daily summarization, a greater amount of storage is needed to retain than daily summarized data.

## Agents that support metric retention summarization

The following agents support summarization for a limited set of metrics: Linux KVM agent, Linux OS agent, UNIX OS agent, and VMware VI agent.

The Cloud App Management console dashboards automatically present the correct visualization of the data based on the available summarized metrics and the time span chosen in the **Events timeline**. Regardless of the retention period that you can choose for hourly and daily summarizations, the metrics that can be summarized cannot be changed.

To get a list of the metrics that can be retained and summarized for each agent, go to developerWorks and download the [ICAM Metric Summarization spreadsheet](#). (See also [Load projection spreadsheet](#) and [Using the IBM Cloud App Management Database Load Projections Spreadsheet](#).)

## Configuring summarization

Summarization is the process of aggregating retained data into time-based categories: hourly and daily. When the metric summary service is enabled, metric samples are summarized and kept on the Cloud App Management server. You can change how often the hourly metrics are summarized, how long hourly and daily summarized metrics are kept on the server, and turn off daily summarization.

### Before you begin

The metric summary service is disabled by default. If you have not enabled summarization, you must run server installation to enable it as described in step “13” on page 147 of *Installing your Cloud App Management server on Red Hat OpenShift*.

### About this task

Use this procedure if you want to change any of these default settings by creating a Kubernetes ConfigMap with the variables that you want to change:

#### **SHORT\_TERM\_SUMMARY\_TTL: "P60D"**

- The hourly summarized metrics for short-term history are kept on the Cloud App Management server for 60 days by default.
- Best practice is to set the value from 8 days to 60 days.

#### **LONG\_TERM\_SUMMARY\_TTL: "P6M"**

- The daily summarized metrics for long-term history are kept on the Cloud App Management server for 6 months by default.
- Best practice is to set the value from 2 months to 13 months.

#### **POLICY\_GENERATION\_HOURLY\_INTERVAL: "PT1H"**

- The metric samples that were taken (typically once per minute) for the past hour are summarized. You can change to another interval in the range 1 - 24 that is evenly divisible into the number of hours in a day: PT1H, PT2H, PT3H, PT4H, PT6H, PT8H, PT12H, or PT24H. If you enter an hourly interval up to 24 that is not one of the previously mentioned values, the interval is rounded down to the nearest divisible value. For example, PT18H is rounded down to PT12H. If you specify a non-integer value, such as PT2.5H, the default PT1H summarization is used.
- The specific hour that hourly summarization is run is determined by adding the hourly interval to midnight Coordinated Universal Time (00:00). For example, summarization for interval PT4H is run at 4:00 AM, 8:00 AM, 12 noon, 4:00 PM, 8:00 PM, and midnight. The exception is PT24H, which indicates that the hourly summarization runs once a day, but the specific hour that it is run is defined by DAILY\_SUMMARIZATION\_HOUR\_START\_TIME because daily summarization cannot run during the first hour of the day.

#### **DAILY\_SUMMARIZATION\_ENABLED: "true"**

- The metric samples that were taken (typically once per minute) for the past day are summarized once per day and saved for long-term history. Creation of the long-term history can take several

hours to complete. If you don't want to keep long-term historical data, use this variable to turn off daily summarization.

- When this variable is used, the **LONG\_TERM\_SUMMARY\_TTL** variable is not used.

### **DAILY\_SUMMARIZATION\_HOUR\_START\_TIME: "01Z"**

- Specifies the hour of the day that daily summarization runs. The default is 1:00 AM UTC. Any hour in the range 1 - 23 can be specified. The value is the UTC hour in the format *hhZ*. For example, to run at 8:00 Eastern Time, you would specify "12Z" because Eastern Time is UTC-4.
- This variable is ignored if **DAILY\_SUMMARIZATION\_ENABLED** is set to "false".
- The daily summarization hour is aligned with the hourly summarization interval, **POLICY\_GENERATION\_HOURLY\_INTERVAL**. If the daily summarization hour is not a value that coincides with the hourly summarization, the daily hour is rounded up to the nearest hourly summarization interval. For example, if daily summarization is set for 2:00 AM and the hourly summarization interval is 4 hours, the daily summarization runs at 4:00 AM. There are two exceptions to this rule:
  - If the hourly summarization runs once a day (**POLICY\_GENERATION\_HOURLY\_INTERVAL**: "PT24H"), both daily and hourly run at the time that is specified for daily summarization. For example, if **POLICY\_GENERATION\_HOURLY\_INTERVAL**: "PT24H" and **DAILY\_SUMMARIZATION\_HOUR\_START\_TIME**: "06Z", then both daily and hourly summarization are run at 06Z (6:00 AM UTC).
  - If rounding up puts the daily summarization at midnight, daily is rounded down to the nearest hourly interval below midnight. For example, if **POLICY\_GENERATION\_HOURLY\_INTERVAL**: "PT4H" and **DAILY\_SUMMARIZATION\_HOUR\_START\_TIME**: "22Z", daily summarization runs at 20Z (8:00 PM UTC).

When you determine the values to set, review the considerations in [“Data retention and summarization”](#) on page 765.

### **Procedure**

Take these steps to configure data summarization:

1. Create the Kubernetes ConfigMap with the following variables, using the ISO 8601 duration format, and save as a YAML file:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: releaseName-metricssummarypolicy-config
  labels:
    release: releaseName
data:
  SHORT_TERM_SUMMARY_TTL: "P60D"
  LONG_TERM_SUMMARY_TTL: "P6M"
  POLICY_GENERATION_HOURLY_INTERVAL: "PT1H"
  DAILY_SUMMARIZATION_ENABLED: "true"
  DAILY_SUMMARIZATION_HOUR_START_TIME: "01Z"
```

where

- *releaseName* is the release name, such as *ibmcloudapmgmt*.
- 60 is the number of days to keep the hourly summarized metrics for short-term history. Best practice is to set the value from 8 days (P8D) to 60 days (P60D).
- 6 is the number of months to keep the daily summarized metrics for long-term history. Best practice is to set the value from 2 months (P2M) to 13 months (P13M).
- 1 is the frequency of hourly summarization. Use this variable to run hourly summarization at a different interval, up to 24 hours (PT24H). For example, set **POLICY\_GENERATION\_HOURLY\_INTERVAL**: PT3H to summarize the past 3 hours of metric samples.

- `true` specifies to perform daily summarization for long-term history, which is set to `true` by default. Set to `false` to disable daily summarization, which also disables the `LONG_TERM_SUMMARY_TTL` variable and causes only hourly summarization to occur.
- `01` is the UTC hour to begin daily summarization.

2. Create the ConfigMap in your Kubernetes environment:

```
kubectl create -f ConfigMap_path
```

where `ConfigMap_path` is the path where you want to save the ConfigMap.

3. Scale the summary policy service down:

```
kubectl scale deployment deployment_name --replicas=0 -n "namespace"
```

where

- `deployment_name` is the name of the deployment of the metric summary policy service. You can find the name with the following command:

```
kubectl get deployments -n namespace | grep metricssummarypolicy
```

- `namespace` is the name of the namespace that was used when you installed Cloud App Management (such as `ops-am`)

4. Scale the summary policy service up:

```
kubectl scale deployment deployment_name --replicas=1 -n "namespace"
```

## Results

The Kubernetes ConfigMap that you created is applied to the summary policy service. If you adjusted the start time, daily summarization occurs on the next day. For example, if daily summarization was originally set to run at 2:00 PM and on Wednesday at 4:00 PM you changed it to run at 11:00 PM, the next summarization occurs on Thursday at 11:00 PM.

# Chapter 20. Monitoring

Use the Cloud App Management console dashboards to manage your resources proactively and effectively troubleshoot issues that arise.

## Viewing your managed resources

Use the **Resources** tab in the Cloud App Management console to get a comprehensive status overview of your microservice-based applications and dynamic workloads that are running in your managed environment. You can drill down to in-depth metrics and adjust the time range to review conditions at the time of an event.

### Before you begin

For any resource that you want to monitor, you must define an incoming event source before you can view data samples and respond to events in the Cloud App Management console. Configure your incoming event source integrations for your Cloud App Management agents or data collectors and any Application Performance Management V8 or IBM Tivoli Monitoring V6 agents that you want to monitor in the Cloud App Management console.

### Procedure

Complete the following steps in the Cloud App Management console to locate and monitor resources across your environment.

1. Click the **Resources** tab.
2. To locate the resources that you want to monitor, complete one of the following steps on the **Resource groups** page.
  - a. Search for a resource group or a resource type: In the Search box, enter any text to search. For example, enter "db". All resource groups and resource types with "db" are shown. Select the resource group or type by clicking the link under the **Name** column.
  - b. **Favorites** pack: From your favorite resource groups pack, select the resource group and drill down to your resources.
  - c. **All resource groups** area: From the list of all resource groups, select the resource group and drill down to your resources.
  - d. **All resource types** area: You can find your resource when you know its resource type. Click the link for your specific resource type under the **Name** column.

A list of available resources and their type are displayed.

3. Drill down the resource metrics for a particular resource by clicking the resource link under the **Resource** column or searching for the resource in the Search box. You can also open the resource dashboard from **Incidents > Events**, click the **See more info** button, and click the link in the URLs section. For more information about the resource dashboard widgets, see ["Resource dashboard"](#) on page 770.

Use the following table to understand the Resource table columns and actions on the **Resources** page.

Resource item	Description
RESOURCE	A link to the <b>Resource</b> dashboard where you can drill down to view the metrics and dashboards that are associated with the resource.
TYPE	The type of resource, for example, LinuxOS. Predefined groups are type: System Defined. A predefined group exists for every type of agent

Resource item	Description
	that you installed in your environment. Custom groups that you or others in your environment define are type: User Defined.
<p>ACTIONS</p> <ul style="list-style-type: none"> <li>• <b>Inspect</b></li> <li>• <b>Thresholds</b></li> <li>• <b>Edit</b></li> </ul>	<p>Open and close actions by clicking the three dot icon on the Resource table. You can inspect a resource, view thresholds, and edit a resource.</p> <ul style="list-style-type: none"> <li>• To drill down into the metrics for a particular resource, click <b>Inspect</b>.</li> <li>• Click the <b>Thresholds</b> icon to view the thresholds that are created for the resource. For example, you might want to view the thresholds if you are monitoring a resource and you notice that a threshold is breached or a resource is encountering an issue. You can drill-down to view the thresholds that are associated with the resource to troubleshoot the issues further. For more information about the thresholds, see <a href="#">“Managing thresholds” on page 755</a>.</li> <li>• Click the <b>Edit</b> icon to edit a resource.</li> </ul>
Resource Count	A drop-down list to display a set number of resources on the <b>Resources</b> page, for example, 15, 25, or 50.

## Resource dashboard

You can use the **Resource** dashboard to monitor your environment by viewing resource metrics across primary resources with the option to drill down further to view the related, secondary resources. When you encounter an issue with your primary resource, you can troubleshoot the related resources to rule out issues at that level.

[Table 1](#) includes a description of the standard monitoring widgets that are displayed for all resource types.

Depending on the resource type, other widgets are shown that are not common across all the dashboards but are specific to a resource. [Table 2](#) includes a description of some of the monitoring widgets that are displayed for other specific resource types such as Kubernetes.

Table 97. Standard monitoring widgets for all resource types

User interface item	Description
<b>Events timeline and time span</b>	<ul style="list-style-type: none"><li>• The default display of the events timeline and chart views presents the past 12 hours. You can adjust the time span to show as few as 3 hours or as much as 1 week. If your Cloud App Management installation is configured for a data retention value of 2 days or more, up to 32 days, the time span options reflect the value that was set. For more information, see <a href="#">“About data retention and summarization” on page 765</a>.</li><li>• For Tivoli Monitoring data providers, a text indicator with the message <code>data provider is online</code> or <code>data provider is offline</code> is displayed on this timeline to show the status of Tivoli Monitoring resource data provider.</li><li>• Square markers are a way to group events that show up in the same vicinity. The number indicates the number of events of the same type that are in close succession. Hover the mouse over an event marker to see when the event was opened and what triggered it. You can click the event to open the corresponding incident.</li><li>• Drag the pin to move across time intervals and view metrics then. For example, if you want to see the metrics at the time an event (or events) occurred, drag the pin to that time.</li></ul>

Table 97. Standard monitoring widgets for all resource types (continued)

User interface item	Description
<p><b>Related Resources</b> widget</p>	<p>View the details and metrics of the secondary or related resources that are associated with your primary resources. Sort any of the following columns in ascending or descending order.</p> <p><b>Search</b> Enter any text to search for a resource, for example, enter LINUX and the LINUX resource is shown.</p> <p><b>Status</b> The status of the resource:</p> <ul style="list-style-type: none"> <li>• Critical</li> <li>• Major</li> <li>• Minor</li> <li>• Warning</li> <li>• Indeterminate</li> <li>• Normal</li> </ul> <p><b>Relation</b> The relationship this related resource has with the one you are monitoring.</p> <p><b>Resource</b> Click the resource link to drill down further to the resource dashboard to look at the metrics more closely.</p> <p><b>Type</b> Predefined groups are type System Defined. You have a predefined group for every type of agent that you installed in your environment. Custom groups that you or others in your environment define are type User Defined.</p>
<p><b>Resource Properties</b> widget</p>	<p>You can view the properties of the object in the monitoring topology service. Scroll through the properties and their values or type in the <b>Filter</b> box to locate a specific property, such as a node's <b>osImage</b> or a pod's <b>qosClass</b>.</p>



Table 97. Standard monitoring widgets for all resource types (continued)


User interface item	Description
<p><b>Custom metrics</b> widget</p>	<p>Click the <b>Custom Metrics</b> twistie and filter the metrics to view the Custom Metrics widgets that are common for all dashboards.</p> <p>Use the <b>Custom Metrics</b> widget to explore the available collected metrics that aren't already reported in the dashboard line charts. You can display up to six additional metrics in one or two line charts:</p> <ol style="list-style-type: none"> <li>1. Click the <b>Custom Metrics</b> &gt; twistie to expand the widget. The widget is typically the last one in the dashboard and shows two views side by side.</li> <li>2. Select an <b>Aggregation</b> function. The aggregations that are available are: Average, Minimum, Maximum, Sum, Deviation.</li> <li>3. Select a metric from the <b>Filter metric</b> list. If the metric has dependencies, a <b>Filter dimension</b> list is displayed. If the <b>Filter dimension</b> has a dependency, another <b>Filter dimension</b> list is displayed or a <b>Dimension value</b> list. The line chart is rendered after you select the required metric and dimension values.</li> <li>4. If a <b>Filter dimension</b> list is displayed, select a dimension from the list. If the metric has multiple dimensions, a second <b>Filter dimension</b> list is displayed for you to select from.</li> <li>5. If a <b>Dimension value</b> list is displayed, select one or more values.</li> </ol> <p>After the required metric type, dimensions, values are selected, the line chart is rendered in the widget space.</p> <p>Example: You might want to view other metrics that relate to the event and correlate these metrics with the standard dashboard metrics. When you view metrics side by side, you can correlate these two sets of metrics.</p>

Depending on the resource type, other widgets are shown that are not common across all the dashboards but they are specific to a resource. For example; if a Linux server has a high CPU usage that caused the incident, you can choose to view a graph that shows the history and trend of CPU utilization, or attributes about the server, or details of the processes that run on the system across various metric widgets.

Table 98. Other monitoring widgets for specific resource types such as Linux or Kubernetes

User interface item	Description
<p><b>Line charts and the golden signals</b></p>	<ul style="list-style-type: none"> <li>Line charts plot metrics from the past three hours or as selected. Hover the mouse pointer over a plot point to see the value and time stamp. All the line charts in the dashboard are synchronized to show the same point in time as you move the mouse pointer across one of the charts.</li> <li>Some of the Kubernetes dashboards have a set of widgets for monitoring the four golden signals: Latency, Errors, Traffic, and Saturation. Latency and Errors typically indicate the symptoms that users are most likely to perceive. The causes behind them are usually Traffic and Saturation.</li> <li>The Latency chart plots the latency in milliseconds. Drag the pin on the timeline or drag the vertical line on the chart to open the hover display for that time. : how long 99% of requests took to complete, how long 50% of requests took to complete, and how long 95% of requests took to complete. For example, a latency of 492 ms in percentile 99 means that 99% of requests took fewer than 492 ms to complete, 189 ms in percentile 50 means that 50% of requests took fewer than 189 ms to complete, and 492 ms in percentile 95 means that 95% of requests took fewer than 492 ms to complete.</li> </ul> <p>Latency</p>  <p>Above the four line charts is the Path widget. If you have multiple end points and only one is performing badly, you can show the signals only for the requests in that path by clearing the check boxes of the other paths.</p>

Table 98. Other monitoring widgets for specific resource types such as Linux or Kubernetes (continued)

User interface item	Description
<p><b>Service Dependencies</b></p>	<p>The Service Dependency view in the Kubernetes Service dashboard shows what application is calling this service and what this service is calling, one step at a time. This view shows service-to-service relationships to help you debug issues across the dependency tree. For example, if the symptoms presented in the Latency and Error line charts are bad but the Traffic and Saturation did not change, you can search this view to find out what is being called. Click a service to open its dashboard.</p> <p>If a Kubernetes service has dependent services, you can click on Expand/Collapse to open a richer topology view. This service dependency view embeds Netcool Agile Service Manager functionality. For more information on using this Service Dependencies, see <a href="#">“Service dependencies topology view”</a> on page 775.</p>
<p><b>Kubernetes topology</b></p>	<ul style="list-style-type: none"> <li>• Hover over the Kubernetes topology to see a pop-up note with status information about an object; click an object to open its dashboard. The topology widget in a <b>Kubernetes Service</b> dashboard displays an  icon if the service provides ingress, and you can click the icon to open the associated Kubernetes Ingress dashboard.</li> <li>• In the Kubernetes Cluster dashboard, click the node to open the associated dashboard. You can drill down to each level in the cluster from the node dashboard to the pod or container dashboard, and from the pod dashboard to the container dashboard. To return to the node dashboard from the pod or container, click inside the hexagon.</li> <li>• From the Kubernetes service, node, pod, or container, you can click the area inside the outermost circle to jump to the cluster dashboard.</li> </ul>

### Service dependencies topology view

To view a topology for Kubernetes services with dependencies, click expand in the Service Dependencies widget. The features of the topology view are described here.

#### Navigation bar

The navigation bar is on the upper left.

#### Number of hops

Select a number between one and four to define the number of relationship hops to be visualized.

#### Type of hops

Choose one of the following hop types:

The **Element to Element** hop type performs the traversal using all element types in the graph.  
The **Host to Host** hop type uses an aggregate traversal across elements with the entity type 'host'.  
The **Element to Host** hop type provides an aggregated hop view like the 'Host to Host' type, but also includes the elements that are used to connect the hosts.

## Visualization toolbar

The Visualization toolbar is available on the left. You can manipulate the topology by using a number of visualization tools.

### Select tool menu

When you hover over the Select tool icon, a submenu is displayed from which you can choose the **Select**, **Pan**, or **Zoom Select** tool.

#### Select tool

Use this icon to select individual resources by using a mouse click, or to select groups of resources by creating a selection area by using click-and-drag.

#### Pan tool

Use this icon to pan across the topology by using click-and-drag on a blank area of the visualization window.

#### Zoom Select tool

Use this icon to zoom in on an area of the topology by using click-and-drag.

### Zoom In

Use this icon to zoom in on the displayed topology.

### Zoom Out

Use this icon to zoom out of the displayed topology.

### Zoom Fit

Use this icon to fit the entire topology in the current view window.

### Overview

Use this icon to create the overview mini map in the lower right corner.

The mini map provides an overview of the entire topology while you zoom in or out of the main topology. The mini map displays a red rectangle to represent the current topology view.

### Layout

Use this icon to recalculate, and then render the topology layout again.

You can choose from a number of layout types and orientations.

#### Layout 1 - Simple topology

A layout that displays all resources in a topology without applying a specific layout structure.

#### Layout 2 - Circular topology

Use when you want to arrange a number of entities by type in a circular pattern.

#### Layout 3 - Grouped topology

Use when you have many linked entities, as it helps you visualize the entities to which a number of other entities are linked. This layout helps to identify groups of interconnected entities and the relationships between them.

#### Layout 4 - Hierarchical topology

Use for topologies that contain hierarchical structures, as it shows how key vertices relate to others with peers in the topology being aligned.

#### Layout 5 = Peacock topology

Use when you have many interlinked vertices, which group the other linked vertices.

#### Layout 6 - system board topology

Use when you want to view how the topology relates to a vertex in terms of its rank, and also how vertices are layered relative to one another.

### Layout 7 - Rank topology

Use when you want to see how a selected vertex and the vertices that are immediately related to it rank relative to the remainder of the topology (up to the specified number of hops). The root selection is automatic.

For example, vertices with high degrees of connectivity outrank lower degrees of connectivity. This layout ranks the topology automatically around the specified seed vertex.

### Layout 8 -Root Rank topology

Similar to a rank topology but it treats the selected vertex as the root. This layout is useful when you want to treat a selected vertex as the root of the tree, with others being ranked below it.

Ranks the topology by using the selected vertex as the root (root selection: Selection)

### Layout orientation

**For layouts 4, 6, 7 and 8**, you can set the following layout orientations:

- Top to bottom
- Bottom to top
- Left to right
- Right to left

### Configure auto update Refresh Rate

Choose 10s, 30s, 1 m, 5 m. When you hover over the **Refresh Rate** icon, a submenu is displayed from which you can configure the auto-update refresh rate.

Click pause auto update to pause topology refresh.

This is unavailable if you are in history mode.

### Open Filter toolbar

The Filter window is displayed on the right, with a **Simple** and **Advanced** tab. Each tab provides you with access to lists of Resource types and Relationship types. Only types relevant to your topology are displayed, for example **host**, **ipaddress** or **operatingsystem**, although you can use the **Show all types** toggle to view all of them.

**Simple tab:** Filter out resource or relationship types, all specified types are removed from view, including the seed resource. By default, all types are **On**. Use the **Off** toggle to remove specific types from your view. It removes **only** the resources that match that type, leaving the resources below, or further out from that type, based on topology traversals.

**Advanced tab:** The Advanced tab performs a server-side topology-based filter action. It removes the resources that match the type, **and** all resources below that type. However, the seed resource is **not** removed from view, even if it is of a type that is selected for removal.

**Reset or invert all filters:** Click **Reset** to switch all types back on, or click **Invert** to invert your selection of types filtered.

**Hover to highlight:** Hover over one of the filtering type options to highlight them in the topology view.

If a filter is applied to a displayed topology, the text 'Filtering applied' is displayed in the status bar at the bottom of the topology.

### Open History toolbar

Use this to open and close the Topology History toolbar. The topology is displayed in history mode by default.

The timeline displays changes to a resource's state, properties, and its relationships with other resources. These changes are displayed through color-coded bars and dash lines, and are elaborated on in a tooltip that is displayed when you hover over the change. You can exclude one or more of these from display.

#### Resource state changes

The timeline displays the state changes for a resource.

#### Resource property changes

The timeline displays the number of times that resource properties were changed.

Each time that property changes were made is displayed as one property change event regardless of whether one or more properties were changed at the time.

### **Resource relationship changes**

The number of relationships with neighboring resources are displayed, and whether these were changed.

The timeline displays when relationships with other resources were changed, and also whether these changes were the removal or addition of a relationship, or the modification of an existing relationship.

To view changes made during a specific time period, use the two time sliders to set the time period. Use the + and - buttons on the right to zoom in and out to increase or decrease the granularity, or by double-clicking within a timeframe. The most granular level you can display is an interval of 1 second. The granularity is depicted with time indicators and parallel bars, which form 'buckets' that contain the recorded resource change event details.

You can use the time picker, which opens a calendar and clock, to move to a specific second in time.

The history timeline is displayed above a secondary time bar, which displays a larger time segment and indicates how much of it is depicted in the main timeline. You can use the jump buttons to move back and forth along the timeline, or jump to the current time.

To view the timeline for a different resource, click it, and the heading above the timeline changes to display the name of the selected resource. If you click the heading, the topology centers (and zooms into) the selected resource.

When you first display the history timeline, coach marks (or tooltips) are displayed, which contain helpful information about the timeline functions. You can scroll through these, or switch them off (or on again) as required.

While in delta mode you can move both pins to show a comparison between the earliest pin and the latest pin. The timeline shows the historic changes for a single selected resource, which is indicated in the timeline title. You can lock one of the time pins in place to be a reference point.

You use the time pins to control the topology shown. When you move the pins, the topology updates to show the topology representation for that time.

### **Context-sensitive view on right-click**

A context-sensitive menu is available when you right-click on a resource.

#### **Menu (right-click)**

Open the menu by using the right-click function. The menu provides access to the following resource-specific actions.

##### **Resource Details**

Displays the current stored properties for the specified resource. Both tabular and raw format are available.

##### **Resource Status**








If statuses related to a specific resource are available, the resource is marked with an icon. The Resource Status option appears in the resource menu.

Displays the time-stamped statuses that are related to the specified resource in table format. The Severity, Time, and State columns can be sorted. The reference time that is shown is the time Resource Status was selected.

If any status tools are defined, the status tool selector (three dots) is displayed next to the resource status. Click the status tool selector to display a list of any status tools that are defined, and then click the specific tool to run it. Status tools are only displayed for the states that were specified when the tools were defined.

The **state** of a status is either 'open', 'clear', or 'closed'.

The **severity** of a status ranges from 'clear' (white tick on a green square) to 'critical' (white cross on a red circle).

Icon	Severity
	Clear
	Indeterminate
	Information
	Warning
	Minor
	Major
	Critical

**Comments**

Displays any comments that are recorded against the resource.

By default, resource comments are displayed by date in ascending order. You can sort them in the following way:

- Oldest first
- Newest first
- User ID (A to Z)
- User ID (Z to A)

Users with the `inasm_operator` role can view comments, but not add any. Users with `inasm_editor` or `inasm_admin` roles can also add new comments.

To add a comment, enter text into the New Comment field, and then click **Add Comment** to save.

**Get Neighbors**

Opens a menu that displays the resource types of all the neighboring resources. Each resource type lists the number of resources of that type and the maximum severity that is associated with each type.

To expand the topology in controlled, incremental steps, choose to get all neighbors of the selected resource, or only the neighbors of a specific type.

Selecting **Get Neighbors** overrides any existing filters.

Click the Undo to return to the previous view.

**Follow Relationship**

Opens a menu that displays all adjacent relationship types.

Each relationship type lists the number of relationships of that type, and the maximum severity that is associated with each type.

You can choose to follow all relationships, or only the neighbors of a specific type.

**Show last change in timeline**

Displays the history timeline, and shows the most recent change that is made to the resource.

**Show first change in timeline**

Displays the history timeline, and shows the first change that is made to the resource.

**Recenter View**

Updates the displayed topology with the specified resource as seed.

**Topology Viewer**

The topology is the central section of the Service dependency view where you view the resource topology

## Resource display conventions

**Deleted:** A minus icon shows that a resource was deleted since last rendered.

Displayed when a topology is updated, and in the history views.

**Added:** A purple plus (+) icon shows that a resource was added since last rendered.

Displayed when a topology is updated, and in the history views.

**Added (neighbors):** A blue asterisk icon shows that a resource was added using the 'get neighbors' function.

## Transaction tracking

The transaction tracking feature enables topology views and instance level transaction monitoring. By distributed tracking infrastructure, transaction tracking can detect bottleneck issues including latency problems and errors, and filter or sort traces based on application. Transaction tracking can also filter views based on length of trace, timestamp, interactions, errors and transaction comparisons.

### About this task

Transaction tracking is installed as part of the IBM Cloud App Management server. Transaction tracking gets the transaction data by OpenTracing. By transaction tracking, you can do the following actions:

- Troubleshoot an issue by looking at a specific request flow through the system.
- Quickly identify the bottleneck which caused the bad experience.
- You can select a single request from the filtered request type shown in the **Service dependencies** topology.
- From the **Service dependencies** topology view, you can expand this instance level topology to view the full topology in context.

The **Service dependencies** view in the **Kubernetes Service** dashboard shows what application is calling this service and what this service is calling, one step at a time. This view shows service-to-service relationships to help you debug issues across the dependency tree. For example, if the symptoms presented in the Latency and Error line charts are bad but the Traffic and Saturation did not change, you can search this view to find out what is being called. Click a service to open its dashboard. For more information about the **Service dependencies** topology, see [“Service dependencies topology view”](#) on page 775.

- View tracing data for the services.

Transaction tracking is automatically enabled for some agents and data collectors but must be manually enabled for others. See the following table for more information about the agents and data collectors that support transaction tracking.

Agent or data collector	Enabled by default	How to enable or disable
Go data collector	✓	<a href="#">“Customizing the Go data collector”</a> on page 582
J2SE data collector	✓	<a href="#">“Customizing the J2SE data collector”</a> on page 588
Liberty data collector	✓	<a href="#">“Customizing the Liberty data collector”</a> on page 597
Node.js data collector	✓	<a href="#">“Customizing the Node.js data collector”</a> on page 604
Python data collector	✓	<a href="#">“Customizing the Python data collector”</a> on page 613



Table 100. Transaction tracking enablement for agents and data collectors (continued)

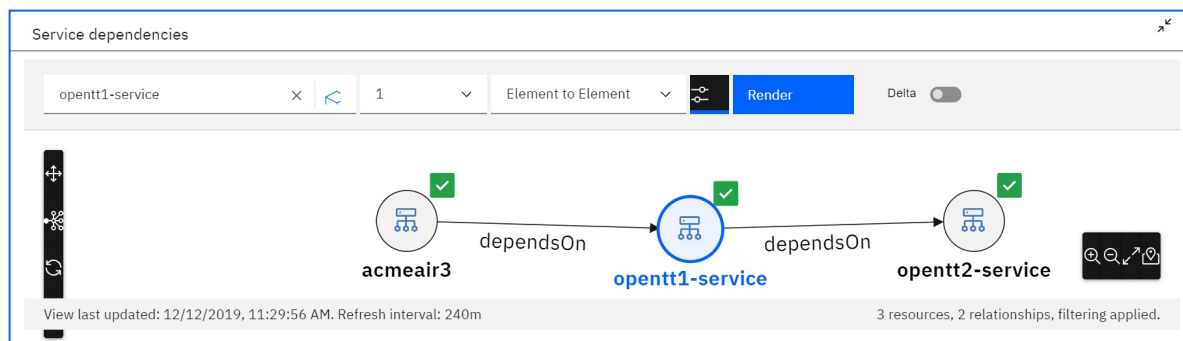
Agent or data collector	Enabled by default	How to enable or disable
Ruby data collector	✓	“Customizing the Ruby data collector” on page 618
WebSphere Applications agent	—	“Configuring the data collector interactively” on page 528
DEM <ul style="list-style-type: none"> <li>• DEM for Liberty applications</li> <li>• DEM for HTTP server</li> </ul>	✓	DEM depends on traction tracking to view browser data and metrics in topology. Do not disable TT for DEM.

**Procedure**

- To view the **Service dependencies** topology, do the following steps:
  - Click the **Resource** dashboard in the UI console.
  - In the **Resource types** list, find **Kubernetes Service** and click it to open.
  - Find the resource name that has transaction tracking enabled, and drill down. You can see the **Service dependencies** topology.




- You can click on Expand/Collapse to open a richer topology view.




- To view the tracing data, do one of the following options:
  - View tracing data from the **Kubernetes Service** dashboard:
    - Click the **Resource** dashboard in the UI console.
    - In the **Resource types** list, find **Kubernetes Service** and click it to open.


3. Find the resource name that has transaction tracking enabled, and drill down.

4. Filter the request in the Golden signals to find the transaction tracing icon  and drill down to check transaction tracing data and topology.

- View tracing data from the runtime data collector resource dashboard:

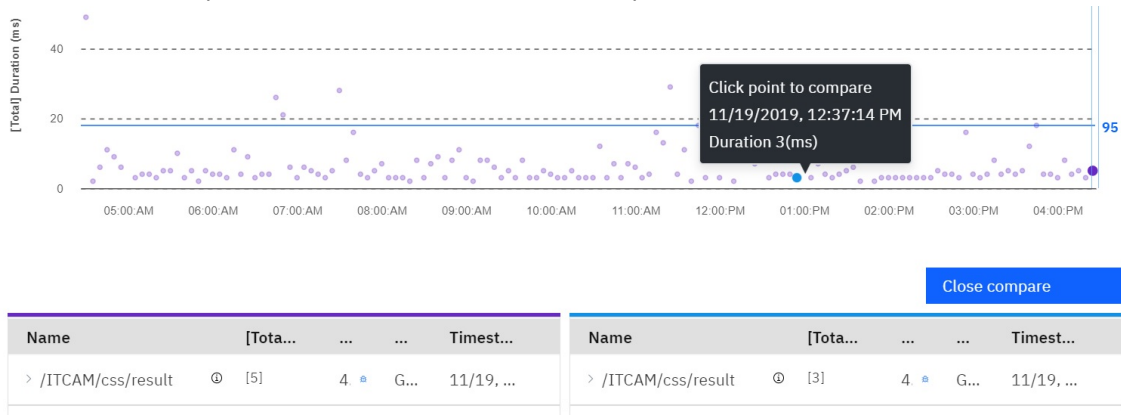
- Go to the Resources tab, and drill down to its dashboard.
- Check the application golden signal data.

3. Filter the request to find the transaction tracing icon  to drill down to check transaction tracing data and topology.



**Note:** For the runtime data collectors, the mechanism for the golden signal sampling rate and the OpenTracing sampling rate are not the same. Therefore, you might not be able to see transaction tracing data for some requests of the golden signal.

- You can drag to move the **5 min** area.
- You can click the point outside the **5 min** area to compare.



## Monitoring the status of your Tivoli Monitoring data providers

If you have issues with your Tivoli Monitoring resources producing monitoring data or you are simply completing regular checks on the health and performance of these resources, you can verify the data

provider status (online or offline) for these specific resources in the event timeline in the **Resource** dashboard.

### Procedure

If one of your resources is not reporting monitoring data in the monitoring dashboards, check the status of the resource data provider by completing the following steps:

1. Click the **Resources** tab in the Cloud App Management console.
2. Under **Resource groups** (either your **Favorites** or **All resource groups**), select the resource group that includes the resources you want to check.
3. Select a specific resource from the list of resources that are displayed.

The events timeline is displayed. A text indicator is displayed on this timeline with the message `data provider is online` or `data provider is offline` is displayed.

If the data provider is offline, the resource is not producing monitoring data so there is no monitoring data being displayed in the charts underneath the timeline. For example, if no data is produced for a 12-hour duration, the monitoring charts are empty. If the data provider is online for some of the 12 hour duration, then subsequently offline and online for the next few hours, you can see a gap in the monitoring data for the time period it was offline.

4. If the data provider is offline, restart it. For more information about restarting your agent data provider, under [Chapter 14, “Configuring the ICAM Agents,”](#) on page 225 information in the Cloud App Management Knowledge Center, go the configuration section for the particular agent you are working with and find the information for starting your agent.

## Viewing Tivoli Monitoring data providers

If you want to quickly view the status and other information about all the Tivoli Monitoring data providers in one view, go to the **Monitoring Data Providers** page in the Cloud App Management console.

### Procedure

View the current list of Tivoli Monitoring data providers by clicking **Monitoring data providers** on the Cloud App Management **Administration** user interface.

A table of Tivoli Monitoring data providers that are sorted in descending order is displayed in the **Monitoring Data Providers** page. The table includes the following information about each data provider:

- The **Name** of the data provider.
- The **Type** of data provider, which is Tivoli Monitoring agent ones currently.
- The **Version** of the data provider.
- The **Hostname** is the data provider host.
- The data provider **Status** is either online or offline. Offline data providers are listed at the top of the table.
- The **Time last changed** is the last time there was a recorded change for the data provider in the topology service.

## ICAM agent and data collector metrics

Your ICAM agents and data collectors monitor the resources in your environment. Each agent and data collector has a set of metrics that are grouped into resource types. Review the metrics topic for your ICAM agent or data collector for a description of the dimensions and metrics in each resource type.

Metrics are the properties that are being measured and reported by your agents and data collectors. Metrics make up the key performance indicators that are presented in the resource page widgets and used to create thresholds for conditions that you want to monitor.

## Db2 agent metrics

The metrics for Db2 agent resource types collect data for monitoring with IBM Cloud App Management. Every Db2 agent resource type defines a set of dimensions and metrics. The descriptions provide such information as data type, dimension key, and metric unit.

### DB2 Database

DB2 Database Instance.

#### Dimensions

##### Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string. This is a key dimension.

##### Instance Name

- The name of the monitored DB2 instance.
- The type is string. This is a key dimension.

##### DB Name

- The real name of the database for which information is collected or to which the application is connected. This name was given to the database when it was created. The value format is a simple text string with a maximum of 60 characters. Use this attribute to identify the specific database to which the data applies.
- The type is string. This is a key dimension.

##### DB Conn Timestamp

- The date and time when the first database connection was made.
- The type is timestamp.

##### Snapshot Timestamp

- The date and time when the database system monitored information was collected. Use this attribute to help correlate data chronologically if you are saving the results in a file or database for ongoing analysis.
- The type is timestamp.

##### DB Path

- The full path of the location where the database is stored on the monitored system. The value format is a simple text string with a maximum of 768 characters. Use this attribute with the Database Name attribute to identify the specific database to which the data applies.
- The type is string.

##### Catalog Node Name

- The network name of the catalog node.
- The type is string.

##### Last Backup

- The date and time that the latest database backup was completed. Use this attribute to help you identify a database that has not been backed up recently, or to identify which database backup file is the most recent. If the database has never been backed up, this timestamp is initialized to zero.
- The type is timestamp.

### Server Platform

- The operating system upon which the database management system is running. Use this attribute during troubleshooting for remote applications.
- The type is string.

### DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If a db partition filter is not specified, data is returned for the current database partition. If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range, the following values are also valid:
- The type is string. This is a key dimension.

### Input DB Alias

- The alias of the database provided when calling the snapshot function. The value format is a simple text string with a maximum of 60 characters. Use this attribute to help you identify the specific database to which the monitor data applies. It contains blanks unless you requested monitor information related to a specific database.
- The type is string.

### Database Status

- The status of the database.
- The type is string.

### DB Location

- The location of the database in relation to the application. Determine the relative location of the database server with respect to the application taking the snapshot.
- The type is string.

## Metrics

### Pool Read Time

- The total amount of elapsed time spent processing read requests that caused data or index pages to be physically read from buffer pool to disk. Use this attribute with the Buffer Pool Data Physical Reads and Buffer Pool Index Physical Reads attributes to calculate the average page-read time. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device. The following value is also valid:
- The type is int.
- The unit is milliseconds.

### Connections Top

- The highest number of simultaneous connections to the database since the database was activated. You can calculate the current number of connections at the time the snapshot was taken by adding the Remote Connections to Database Manager and Local Connections attributes. Use this attribute to evaluate the setting of the MAXAPPLS configuration parameter. The following value is also valid:
- The type is int.
- The unit is connections.

### Select SQL Stmt

- The number of SQL SELECT statements that ran. Use this attribute to determine the level of database activity at the application or database level. You can also use the following formula to determine the ratio of SELECT statements to the total statements by performing the following operations:

- Add the number of static SQL statements attempted and dynamic SQL statements attempted
- Divide the resulting total by the number of select SQL statements that ran

The following value is valid:

- The type is int.
- The unit is selects.

#### Pool Data Writes

- The number of times a buffer pool data page was physically written to disk. A buffer pool data page is written to disk for the following reasons:
  - To free a page in the buffer pool so another page can be read
  - To flush the buffer pool.

If a buffer pool data page is written to disk for a high percentage of Buffer Pool Data Physical Reads, performance might improve by increasing the number of buffer pool pages available for the database.

- The type is int.
- The unit is writes.

#### Pool LSN Gap Clns

- The number of times a page cleaner was invoked because the logging space used had reached a predefined criterion for the database. Use this attribute to evaluate whether you have enough space for logging, and whether you need more log files or larger log files. The following value is also valid:
- The type is int.
- The unit is invocations.

#### Dynamic SQL Stmtns

- The number of dynamic SQL statements that were attempted. Use this attribute to calculate the total number of successful SQL statements at the database or application level by performing the following operations:

1. Add the number of Dynamic SQL Statements Attempted and the Static SQL Statements Attempted.
2. Subtract the number of Failed Statement Operations.

The remainder equals the throughput (the number of successful SQL statements) during the current monitoring period. The following value is also valid:

- The type is int.
- The unit is statements.

#### Num Assoc Agents

- The current number of subagents associated with all applications that are connected to the monitored database. Use this attribute to evaluate the settings for the agent configuration parameters. The following value is also valid:
- The type is int.
- The unit is subagents.

#### Direct Write Time

- The elapsed time (in milliseconds) required to perform the direct writes. Use the following formula to calculate the average direct write time per sector: direct write time / direct writes to database A high average time might indicate an I/O conflict. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### DDL SQL Stmts

- The number of SQL Data Definition Language (DDL) statements that ran. Use this attribute to determine the level of database activity at the application or database level. DDL statements are expensive to run because of their impact on the system catalog tables. As a result, if the value of this attribute is high, you must determine the cause and possibly restrict the identified activity from being performed. You can also use this attribute to determine the percentage of DDL activity using the following formula: divide the number of DDL SQL statements by the total number of statements. The following value is also valid:
- The type is int.
- The unit is statements.

#### Longest Lock Wait Time

- The longest time for which an application waited for a lock.
- The type is int.
- The unit is milliseconds.

#### Int Rollbacks

- The total number of rollbacks initiated internally by the database manager. An internal rollback occurs when any of the following operations cannot complete successfully:
  - A reorganization
  - An import
  - A bind or pre-compile
  - An application that ends as a result of a deadlock situation or lock timeout situation
  - An application that ends without executing an explicit COMMIT or ROLLBACK statement (on Windows systems).

Use this attribute to calculate the total number of units of work by calculating the sum of the following values: commit statements attempted, internal commits, rollback statements attempted, and internal rollbacks.

- The type is int.
- The unit is rollbacks.

#### Pool Async Data Reads

- The number of pages read asynchronously into the buffer pool. Use this attribute with the Buffer Pool Data Physical Reads attribute to calculate the number of physical reads that were performed synchronously (that is, physical data page reads that were performed by database manager agents). Use the following formula: buffer pool data physical reads - buffer pool synchronous data reads. By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working.
- The type is int.
- The unit is pages.

#### Direct Read Reqs

- The number of requests to perform a direct read of one or more sectors of data. Use the following formula to calculate the average number of sectors that are read by a direct read:  $\text{direct reads from database} / \text{direct read requests}$  The following value is also valid:
- The type is int.
- The unit is requests.

#### Files Closed

- The total number of database files closed. The database manager opens files for reading and writing into and out of the buffer pool. The maximum number of database files open by an application at any time is controlled by the MAXFILOP configuration parameter. If the maximum is reached, one file is closed before the new file is opened. Note that the actual number of files opened might not equal the number of files closed. The following value is also valid:
- The type is int.
- The unit is files.

#### Total Hash Joins

- The total number of hash joins that ran. At the database or application level, use this value with the Hash Join Overflows attribute and the Hash Join Small Overflows attribute to determine if a significant percentage of hash joins would benefit from modest increases in the sort heap size. The following value is valid:
- The type is int.
- The unit is joins.

#### Total Cons

- The number of connections to the database since the first connect, activate, or last reset (coordinator agents). Use this attribute with the Database Activation Timestamp and the Start Database Manager Timestamp attributes to calculate the frequency at which applications have connected to the database. The first connect to a database (such as an initial buffer pool allocation) causes extra overhead. If the frequency of connects is low, it might be beneficial to activate the database explicitly using the ACTIVATE DATABASE command before connecting any other application. As a result, subsequent connects are processed at a higher rate. The following value is valid: When you reset this attribute, the value of the attribute is set to the number of applications that are currently connected, instead of zero.
- The type is int.
- The unit is connections.

#### Failed SQL Stmts

- The number of SQL statements that were attempted, but failed. This count includes all SQL statements that received a negative SQLCODE. Use this attribute to calculate the total number of successful SQL statements at the database or application level by performing the following operations:

1. Add the number of Dynamic SQL Statements Attempted and the Static SQL Statements Attempted.
2. Subtract the number of Failed Statement Operations.

The remainder equals the throughput (the number of successful SQL statements) during the current monitoring period. This attribute can also help you to determine the reasons for poor performance; failed statements indicate time wasted by the database manager, which results in lower throughput for the database.

- The type is int.
- The unit is failures.



#### Total Log Used Pct

- The percentage of used log space in the database.
- The type is double.
- The unit is percent.

#### Lock Escals

- The number of times the lock escalations occurred for several row locks to a table lock. A lock is escalated when the total number of locks that are held by an application reaches the maximum amount of lock list space available to the application, or the lock list space that the applications use approaches the total lock list space. This data item includes the count of all lock escalations, including exclusive lock escalations. When an application reaches the maximum number of locks that are permitted and no locks are available to escalate, the application uses the space in the lock list that is allocated for other applications. When the entire lock list is full, an error occurs. The value format is an integer.
- The type is int.
- The unit is escalations.

#### Lock Waits

- The total number of times when the applications or connections waited for locks. At the database level, the lock waiting value is the total number of times that applications waited for locks within the database. At the application-connection level, the lock waiting value is the total number of times that this connection requested a lock but waited because another connection was already holding a lock on the data. Use this attribute to calculate the average wait time for a lock. If the average lock wait time is high, you must look for applications that hold many locks, or have lock escalations, with a focus on tuning your applications to improve concurrency.
- The type is int.
- The unit is waits.

#### SQL Stmts Failed Percent

- The percentage of SQL statements that failed to run successfully. This value is derived by dividing the value of the Failed SQL Statements attribute by the value of the Total SQL Statements attribute. Use this attribute to determine whether an application has some design issues.
- The type is double.
- The unit is percent.

#### Total Log Used Top

- The maximum amount of total log space (in bytes) that has been used. Use this attribute to evaluate the amount of primary log space that is allocated. Comparing the value of this attribute with the amount of primary log space that is allocated can help you to evaluate the configuration parameter settings. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal. The following value is valid:
- The type is int.
- The unit is bytes.

#### Pool Total Writes

- The total number of write requests. This attribute is the total of the Pool Data Writes and Pool Index Writes attributes. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal.
- The type is int.
- The unit is writes.

### Avg Pool Read Time

- The average elapsed time for a read request. This value is derived by dividing the value of the Pool Read Time attribute by the value of the Pool Total Reads attribute. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device. The following value is also valid:
- The type is int.
- The unit is milliseconds.

### Int Deadlock Rollbacks

- The total number of forced rollbacks initiated by the database manager due to a deadlock. The database manager initiates a rollback for the current unit of work in an application that is experiencing a deadlock. This attribute shows the number of deadlocks that have been broken. It can indicate the possibility of concurrency problems. It is also important because internal rollbacks due to deadlocks can cause performance degradation. The following value is also valid:
- The type is int.
- The unit is rollbacks.

### Pool Index Writes

- The number of times a buffer pool index page was physically written to the disk. If this number is high, increase the number of buffer pool pages for the database to improve the performance.
- The type is int.
- The unit is writes.

### Pool Async Index Writes

- The number of times a buffer pool index page was physically written to disk by an asynchronous page cleaner or a prefetcher. A prefetcher might have written dirty pages to disk to make space for the pages being prefetched. Use this attribute with the Buffer Pool Index Writes attribute to calculate the number of physical index write requests that were performed synchronously. That is, physical index page writes that were performed by database manager agents. Use the following formula:  $\text{buffer pool index writes} - \text{buffer pool asynchronous index writes}$  By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing.
- The type is int.
- The unit is writes.

### Pool Async Data Writes

- The number of times a buffer pool data page was physically written to disk by an asynchronous page cleaner or by a prefetcher. A prefetcher might have written dirty pages to disk to make space for the pages being prefetched. Use this attribute with the Buffer Pool Data Writes attribute to calculate the number of physical write requests that were performed synchronously (that is, physical data page writes that were performed by database manager agents). Use the following formula:  $\text{buffer pool data writes} - \text{buffer pool asynchronous data writes}$  By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing.
- The type is int.
- The unit is writes.

### Static SQL Stmts

- The number of static SQL statements that were attempted. Use this attribute to calculate the total number of successful SQL statements at the database or application level by performing the following operations:

1. Add the number of Dynamic SQL Statements Attempted and the Static SQL Statements Attempted.
2. Subtract the number of Failed Statement Operations.

The remainder equals the throughput (the number of successful SQL statements) during the current monitoring period.

- The type is int.
- The unit is statements.

#### Sort Heap Allocated

- The total number of allocated pages of sort heap space for all sorts at the level chosen (database manager or database) and at the time the snapshot was taken. Memory estimates do not include sort heap space. If excessive sorting occurs, add the extra memory (used for the sort heap) to the base memory requirements for running the database manager. Generally, the larger the sort heap, the more efficient the sort. Appropriate use of indexes can reduce the amount of sorting required.
- The type is int.
- The unit is pages.

#### Maximum Connection

- The maximum number of applications that can connect to the monitored database.
- The type is int.
- The unit is applications.

#### Sec Log Used Top

- The maximum amount of secondary log space (in bytes) that has been used. Use this attribute with the Secondary Logs Allocated and Total Log Used Top attributes to show the current dependency on secondary logs. If this value is high, you might need larger log files, more primary log files, or more frequent COMMIT statements within your application. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal. The following value is valid:
- The type is int.
- The unit is bytes.

#### Direct Reads

- The number of read operations that occurred without using the buffer pool. Use the following formula to calculate the average number of sectors that are read by a direct read: direct reads from database/direct read requests. When you use the system monitors to track I/O, this attribute helps to distinguish a database I/O from a non-database I/O on the device.
- The type is int.
- The unit is reads.

#### Rows Updated

- The number of row updates attempted. Use this attribute to gain insight into the current level of activity within the database manager. The following value is also valid:
- The type is int.
- The unit is updates.

#### Pool Index from Estore

- Number of buffer pool index pages copied from extended storage. Required index pages are copied from extended storage to the buffer pool. The copy process might incur the cost of connecting to the shared memory segment, but it saves the cost of a disk read. The following value is also valid:

- The type is int.
- The unit is pages.

#### Rows Deleted

- The number of row deletions attempted. Use this attribute to gain insight into the current level of activity within the database manager. The following value is also valid:
- The type is int.
- The unit is deletes.

#### Lock Timeouts

- The number of times when a request to lock an object timed out instead of being granted. Use this attribute to adjust the setting for the LOCKTIMEOUT database configuration parameter. The high lock timeouts indicate that an application might be holding locks for long duration. In such a case, analyze some other attributes related to locks and deadlocks to determine whether the application problem exists. If the value for the LOCKTIMEOUT parameter is set too high, applications might wait for longer duration to acquire a lock.
- The type is int.
- The unit is timeouts.

#### Total Sort Time

- The total elapsed time (in milliseconds) for all sorts that ran. at the database or application level, use this attribute with the Total Sorts attribute to calculate the average sort time, which can indicate whether sorting is a performance issue. Elapsed times are affected by system load. The more processes you have running, the higher this elapsed time value is. The following value is valid:
- The type is int.
- The unit is milliseconds.

#### Total Sec Cons

- The number of connections made by a subagent to the database at the node. Use this attribute with the Connects Since Database Activation, Database Activation Timestamp, and the Start Database Manager Timestamp attributes to calculate the frequency at which applications have connected to the database. The following value is valid:
- The type is int.
- The unit is connections.

#### Sec Logs Allocated

- The total number of secondary log files that are currently being used for the database. Use this attribute with the Secondary Log Used Top and Total Log Used Top attributes to show the current dependency on secondary logs. If this value is consistently high, you might need larger log files, more primary log files, or more frequent COMMIT statements within your application. The following value is valid:
- The type is int.
- The unit is files.

#### Hash Join Small Overflows

- The number of times that hash join data exceeded the available sort heap space by less than 10%. If this value and the value of the Hash Join Overflows attribute are high, consider increasing the sort heap threshold. If this value is greater than 10% of Hash Join Overflows, consider increasing the sort heap size. The following value is also valid:
- The type is int.

- The unit is overflows.

#### Direct Write Reqs

- The number of requests to perform a direct write of one or more sectors of data. Use the following formula to calculate the average number of sectors that are written by a direct write:  $\text{direct writes to database} / \text{direct write requests}$ . The following value is also valid:
- The type is int.
- The unit is requests.

#### Transaction Per Min

- The total number of database transactions that occurred per minute.
- The type is int.
- The unit is transactions/minute.

#### Deadlocks

- The total number of deadlocks that have occurred. This attribute can indicate that applications are experiencing contention problems. To resolve the problem, determine in which applications (or application processes) the deadlock are occurring. You can then modify the application to enable it to run concurrently. Some applications, however, might not be capable of running concurrently.
- The type is int.
- The unit is deadlocks.

#### Rows Inserted

- The number of row insertions attempted. Use this attribute to gain insight into the current level of activity within the database manager. The following value is also valid:
- The type is int.
- The unit is inserts.

#### Pool Drty Pg Steal Clns

- Buffer Pool Victim Page Cleaners Triggered is the number of times a page cleaner was invoked because a synchronous write was needed during the victim buffer replacement for the database. Use this attribute, in combination with others, to evaluate the number of page cleaners that are defined.
- The type is int.
- The unit is invocations.

#### Pkg Cache Hit Ratio

- The percentage of package sections that were found in the cache. This ratio tells you whether the package cache is being used effectively. If the hit ratio is high (more than 0.8), the cache is performing well. A smaller ratio might indicate that the package cache must be increased.
- The type is double.
- The unit is percent.

#### Pool Data L Reads

- The number of logical read requests for data pages that have gone through the buffer pool. This count includes accesses to the following data:
  - Data that is already in the buffer pool when the database manager needs to process the page
  - Data that is read into the buffer pool before the database manager can process the page.

By using the Pool Data Physical Reads attribute, you can calculate the data page hit ratio for the buffer pool as follows:  $1 - (\text{buffer pool data physical reads} / \text{buffer pool data logical reads})$  By using the Pool Data Physical Reads, Pool Index Physical Reads, and Pool Index Logical Reads attributes, you can calculate the overall buffer pool hit ratio as follows:  $1 - ((\text{buffer pool data physical reads} + \text{buffer pool index physical reads}) / (\text{buffer pool data logical reads} + \text{buffer pool index logical reads}))$  Increasing buffer pool size generally improves the hit ratio until you reach a point of diminishing returns.

- The type is int.
- The unit is requests.

#### Avg Lock Wait Time

- The average elapsed time (in milliseconds) that was spent waiting for a lock. If the average lock wait time is high, identify the applications that hold many locks, or have lock escalations, with a focus on tuning your applications to improve concurrency, if appropriate. If the average lock waiting time is high due to escalations, the values of one or both of the LOCKLIST and MAXLOCKS configuration parameters might be too low.
- The type is int.
- The unit is milliseconds.

#### Database Application Locks Held

- The number of locks that the applications currently hold. If the monitor information is at the database level, this number represents the total number of locks the database applications currently hold. If the information is at the application level, this number represents the total number of locks the application agents currently hold.
- The type is int.
- The unit is locks.

#### Pool Async Index Reads

- The number of index pages read asynchronously into the buffer pool by a prefetcher. Asynchronous reads are performed by database manager prefetchers. Use this attribute with the Buffer Pool Index Physical Reads attribute to calculate the number of physical reads that were performed synchronously (that is, physical index page reads that were performed by database manager agents). Use the following formula:  $\text{buffer pool index physical reads} - \text{buffer pool asynchronous index reads}$  By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working.
- The type is int.
- The unit is reads.

#### Hash Join Overflows

- The number of times that hash join data exceeded the available sort heap space. At the database level, if the percentage of Hash Join Small Overflows is greater than 10% of this value, you must consider increasing the sort heap size. You can use values at the application level to evaluate hash join performance for individual applications. The following value is also valid:
- The type is int.
- The unit is overflows.

#### Pool Index to Estore

- Number of buffer pool index pages copied to extended storage. Pages are copied from the buffer pool to extended storage when they are selected as victim pages. As a result of the copying process, there is sufficient space for new pages in the buffer pool. The following value is also valid:
- The type is int.

- The unit is pages.

#### Cat Cache Overflows

- The number of times that an insert into the catalog cache failed because the catalog cache was full. If the catalog cache overflows value is large, the catalog cache might be too small for the workload. Increasing the size of the catalog cache might improve its performance. If the workload includes transactions that compile a large number of SQL statements referencing many tables, views, and aliases in a single unit of work, compiling fewer SQL statements in a single transaction might improve the performance of the catalog cache. Or if the workload includes the binding of packages containing many SQL statements referencing many tables, views or aliases, you might want to split the packages so that they include fewer SQL statements to improve performance. The following value is also valid:
  - The type is int.
  - The unit is failures.

#### Locks Waiting

- The number of agents that are currently waiting on a lock.
  - The type is int.
  - The unit is agents.

#### Active Sorts

- The number of sorts in the database with the allocated sort heap. Use this value with the Sort Heap Allocated attribute to determine the average sort heap space that is used by each sort. If the SORTHEAP configuration parameter is substantially larger than the average sort heap used, you can reduce the value of this parameter.
  - The type is int.
  - The unit is sorts.

#### Direct Writes

- The number of write operations that occurred without using the buffer pool. Use the following formula to calculate the average number of sectors that are written by a direct write: direct writes to database/direct write requests. When you use the system monitors to track I/O, this attribute helps to distinguish a database I/O from a non-database I/O on the device.
  - The type is int.
  - The unit is writes.

#### X Lock Escals

- The number of times that locks have been escalated from several row locks to one exclusive table lock, or the number of times an exclusive lock on a row caused the table lock to become an exclusive lock. A lock is escalated when the total number of locks held by an application reaches the maximum amount of lock list space available to the application. The amount of lock list space available is determined by the LOCKLIST and MAXLOCKS configuration parameters. Other applications cannot access data held by an exclusive lock. Because exclusive locks can affect the concurrency of your data, it is important to track them. When an application reaches the maximum number of locks allowed and there are no more locks to escalate, it uses space in the lock list allocated for other applications. When the entire lock list is full, an error occurs. See the Lock Escals attribute for possible causes and resolutions to excessive exclusive lock escalations. The following value is valid:
  - The type is int.
  - The unit is escalations.

#### Lock List in Use

- The total amount of lock list memory (in bytes) that is currently in use. This attribute can be used with the locklist configuration parameter to calculate the lock list utilization. If the lock list utilization is high, you might want to consider increasing the size of that parameter. Values that are greater than or equal to 2147483647 are indicated in the portal with the Value Exceeds Maximum text, and values that are smaller than -2147483648 are indicated with the Value Exceeds Minimum text. The following value is also valid:
- The type is int.
- The unit is bytes.

#### Pool Hit Ratio

- The buffer pool hit ratio (as a percentage). The sum of the Pool Data Logical Reads and Pool Index Logical Reads attributes is divided by the value of the Pool Total Reads attribute to derive the pool hit ratio. Use this attribute to determine whether buffer pool assignment is efficient. If the pool hit ratio is low, increasing the number of buffer pool pages might improve performance.
- The type is double.
- The unit is percent.

#### Pool Write Time

- The total amount of time spent physically writing data or index pages from the buffer pool to disk. Use this attribute with the Buffer Pool Data Writes and Buffer Pool Index Writes attributes to calculate the average page-write time. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Avg Sort Time

- The average derived by dividing value of the Total Sort Time attribute by the value of the Total Sorts attribute. The average is expressed as elapsed time. at the database or application level, this attribute can indicate whether sorting is a performance issue. Elapsed times are affected by system load. The more processes you have running, the higher this elapsed time value is.
- The type is int.
- The unit is Milliseconds.

#### Pool Data to Estore

- Number of buffer pool data pages copied to extended storage. Pages are copied from the buffer pool to extended storage when they are selected as victim pages. As a result of the copying process, there is sufficient space for new pages in the buffer pool. The following value is also valid:
- The type is int.
- The unit is pages.

#### Total SQL Stmts

- The total number of dynamic and static SQL statements. This value is derived by adding the values of the Dynamic SQL Statements and the Static SQL Statements attributes. The following value is valid:
- The type is int.
- The unit is statements.

#### SQL Stmts Rollback Percent



- The percentage of SQL statements that resulted in a rollback. This value is derived by dividing the value of the Rollback SQL Statements attribute by the value of the Total SQL Statements attribute. Use this attribute to determine whether an application has some design issues.
- The type is double.
- The unit is percent.

#### Pool Sync Data Reads

- The number of physical data page reads that were performed by database manager agents. This value is derived by subtracting the value of the Pool Async Data Reads attribute from the Pool Data Physical Reads attribute. By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working. The following value is also valid:
- The type is int.
- The unit is reads.

#### Rows Selected

- The number of rows that have been selected and returned to the application. Use this attribute to gain insight into the current level of activity within the database manager. The following value is also valid:
- The type is int.
- The unit is selects.

#### Pool Total Reads

- The total number of read requests that required I/O to get data pages and index pages into the buffer pool. This attribute is the total of the Pool Data Physical Reads and Pool Index Physical Reads attributes. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal.
- The type is int.
- The unit is reads.

#### Pkg Cache Inserts

- The total number of times that a requested section was not available for use and had to be loaded into the package cache. This count includes any implicit prepares performed by the system. By using the Package Cache Lookups attribute, you can calculate the package cache hit ratio using the following formula:  $1 - (\text{Package Cache Inserts} / \text{Package Cache Lookups})$  See the Package Cache Lookups attribute for information about using this attribute.
- The type is int.
- The unit is occurrences.

#### Pool Sync Index Reads

- The number of index pages read synchronously (that is, physical index page reads that were performed by database manager agents) into the buffer pool. This value is derived by subtracting the value of the Pool Async Index Reads attribute from Pool Index Physical Reads attribute. By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working. The following value is also valid:
- The type is int.
- The unit is reads.

#### Pool Index P Reads

- The number of physical read requests to get index pages into the buffer pool. See the Pool Index Logical Reads attribute for information about how to use this element.

- The type is int.
- The unit is reads.

#### Sort Overflows

- The total number of sorts that ran out of sort heap space and might have required disk space for temporary storage. at the database or application level, use this attribute with the Total Sorts attribute to calculate the percentage of sorts that required overflow to disk. If this percentage is high, you might want to adjust the database configuration by increasing the value of the SORTHEAP configuration parameter. The following value is valid:
- The type is int.
- The unit is sorts.

#### Sort Overflows Percent

- The percentage of sorts that ran out of sort heap space and might have required disk space for temporary storage. This percentage is calculated by dividing the value of the Sort Overflows attribute by the value of the Total Sorts attribute. at the database or application level, use this attribute to evaluate the percentage of sorts that required overflow to disk. If this percentage is high, you might want to adjust the database configuration by increasing the value of the SORTHEAP configuration parameter.
- The type is double.
- The unit is percent.

#### Memory Used Pct

- The percentage usage of system memory by the database.
- The type is double.
- The unit is percent.

#### Pool Drty Pg Thrsh Clns

- The number of times a page cleaner was invoked because a buffer pool had reached the dirty page threshold criterion for the database. When the number of dirty pages in the pool exceeds this value, the cleaners are triggered. If this value is set too low, pages might be written out too early, requiring them to be read back in. If set too high, too many pages might accumulate, requiring users to write out pages synchronously.
- The type is int.
- The unit is invocations.

#### Total Hash Loops

- The total number of times that a single partition of a hash join was larger than the available sort heap space. This attribute might indicate inefficient execution of hash joins (the sort heap size is too small or the sort heap threshold is too small). Use this value with the other hash join variables to tune the sort heap size (SORTHEAP) and sort heap threshold (SHEAPTHRES) configuration parameters. The following value is valid:
- The type is int.
- The unit is occurrences.

#### Lock Wait Time

- The total elapsed time (in milliseconds) that was spent waiting for a lock. At the database level, this is the total amount of elapsed time of all applications that were waiting for a lock within this database. At the application-connection and transaction levels, this is the total amount of elapsed time that this connection or transaction waited for a lock to be granted. Use this attribute with the Lock Waits attribute to calculate the average wait time for a lock. This calculation can be performed

at the database or the application connection level. If the average wait time for a lock is high, identify the applications with many locks or lock escalations for improving concurrency of applications. If the average wait time for a lock is high due to escalations, the values of one or both of the LOCKLIST and MAXLOCKS configuration parameters might be too low.

- The type is int.
- The unit is milliseconds.

#### Commit SQL Stmts

- The total number of SQL COMMIT statements that have been attempted. A small rate of change in this counter during the monitor period might indicate that applications are not doing frequent commits. The lack of frequent commits can lead to problems with logging and data concurrency. You can also use this attribute to calculate the total number of units of work by calculating the sum of the following values:

- Commit statements attempted
- Internal commits
- Rollback statements attempted
- Internal rollbacks

The following value is also valid:

- The type is int.
- The unit is commits.

#### Cat Cache Inserts

- The number of times that the system tried to insert table descriptor information into the catalog cache. Table descriptor information is inserted into the cache following a failed lookup to the catalog cache while processing a table, view, or alias reference in an SQL statement. The catalog cache inserts value includes attempts to insert table descriptor information that fail due to catalog cache overflow and heap full conditions. The following value is also valid:
- The type is int.
- The unit is attempts.

#### Avg Pool Write Time

- The average elapsed time for a write request. This value is derived by dividing the value of the Pool Write Time attribute by the value of the Pool Total Writes attribute. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Direct Read Time

- The elapsed time (in milliseconds) required to perform the direct reads. Use the following formula to calculate the average direct read time per sector:  $\text{direct read time} / \text{direct reads from database A}$ . A high average time might indicate an I/O conflict. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Log Writes

- The number of log pages written to disk by the logger. Use this attribute with an operating system monitor to quantify the amount of I/O on a device that is attributable to database activity. The following value is also valid:
- The type is int.
- The unit is pages.

### Cat Cache Lookups

- The number of times that the catalog cache was referenced to obtain table descriptor information. This attribute includes both successful and unsuccessful accesses to the catalog cache. This attribute is used in calculating the catalog cache hit ratio. This ratio indicates how well the catalog cache is avoiding catalog accesses. If the ratio is high (more than 0.8), the cache is performing well. A smaller ratio might indicate that you must increase the size of the catalog cache. You must expect a large ratio immediately following the first connection to the database. The following value is also valid:
  - The type is int.
  - The unit is lookups.

### Coord Agents Top

- The maximum number of coordinating agents working at one time. The MAXCAGENTS configuration parameter determines the number of coordinating agents that can be executing concurrently. If the peak number of coordinating agents results in a workload that is too high for this node, you can reduce the MAXCAGENTS configuration parameter. The following value is also valid:
  - The type is int.
  - The unit is agents.

### Agents Top

- The maximum number of agents (at one time) associated with applications that are connected to the monitored database. Use this attribute to indicate how well the intra-query parallelism was realized. The following value is also valid:
  - The type is int.
  - The unit is agents.

### Rollback SQL Stmts

- The total number of SQL ROLLBACK statements that have been attempted. A rollback can result from an application request, a deadlock, or an error situation. This attribute counts only the number of rollback statements issued from applications. At the application level, this attribute can help you determine the level of database activity for the application and the amount of conflict with other applications. At the database level, it can help you determine the amount of activity in the database and the amount of conflict between applications on the database. The following value is also valid:
  - The type is int.
  - The unit is rollbacks.

### Pool Async Write Time

- The total elapsed time spent writing data or index pages from the buffer pool to disk by database manager page cleaners. Calculate the elapsed time spent writing pages synchronously by subtracting the value of the Pool Async Write Time attribute from the value of the Pool Physical Write Time attribute. You can also use this attribute to calculate the average asynchronous read time by performing the following operations:
  1. Add the Pool Async Data Writes and the Pool Async Index Writes.
  2. Divide the Pool Async Write Time by the sum from step 1.These calculations can be used to understand the I/O work being performed.
  - The type is int.
  - The unit is milliseconds.

### Pool Data P Reads

- The number of read requests that required I/O to get data pages into the buffer pool. See Pool Data Logical Reads and Pool Async Data Reads attributes for information about how to use this attribute. The following value is also valid:
- The type is int.
- The unit is requests.

#### Pool Index L Reads

- The number of logical read requests for index pages that have gone through the buffer pool. This count includes accesses to the following index pages:
  - Pages that are already in the buffer pool when the database manager needs to process the page.
  - Pages that are read into the buffer pool before the database manager can process the page.

By using the Buffer Pool Index Physical Reads attribute, you can calculate the index page hit ratio for the buffer pool as follows:  $1 - (\text{buffer pool index physical reads} / \text{buffer pool index logical reads})$ . To calculate the overall buffer pool hit ratio, see the Buffer Pool Data Logical Reads attribute. If the hit ratio is low, increasing the number of buffer pool pages might improve performance.

- The type is int.
- The unit is requests.

#### Appls Cur Cons

- The number of applications that are currently connected to the monitored database. Use this attribute to understand the level of activity within a database and the amount of system resource that is being used.
- The type is int.
- The unit is applications.

#### Pool Data from Estore

- Number of buffer pool data pages copied from extended storage. Required pages are copied from extended storage to the buffer pool. The copy process might incur the cost of connecting to the shared memory segment, but it saves the cost of a disk read.
- The type is int.
- The unit is pages.

#### CPU Used

- The percentage usage of CPU by the database.
- The type is int.
- The unit is milliseconds.

#### Pkg Cache Lookups

- The number of times that an application looked for a section or package in the package cache. At a database level, it indicates the overall number of references since the database was started, or monitor data was reset. Note that this counter includes the cases where the section is already loaded in the cache and when the section has to be loaded into the cache. To calculate the package cache hit ratio use the following formula:  $1 - (\text{Package Cache Inserts} / \text{Package Cache Lookups})$ . The package cache hit ratio tells you whether the package cache is being used effectively. If the hit ratio is high (more than 0.8), the cache is performing well. A smaller ratio might indicate that the package cache must be increased.
- The type is int.
- The unit is lookups.

#### Cat Cache Heap Full

- The number of times that an insert into the catalog cache failed because of a heap full condition in the database heap. The catalog cache draws its storage dynamically from the database heap. Even if the cache storage has not reached its limit, inserts into the catalog cache might fail due to a lack of space in the database heap. If the catalog cache heap full count is not zero, you can correct the insert failure condition by increasing the database heap size or by reducing the catalog cache size. The following value is also valid:
- The type is int.
- The unit is failures.

#### Log Reads

- The number of log pages read from disk by the logger. The following value is also valid:
- The type is int.
- The unit is pages.

#### Cat Cache Hit Ratio

- The percentage of catalog sections in the cache. This ratio indicates how well the catalog cache is avoiding catalog accesses. If the ratio is high (more than 0.8), the cache is performing well. A smaller ratio might indicate that you must increase the size of the catalog cache. You must expect a large ratio immediately after the first connection to the database.
- The type is double.
- The unit is percent.

#### Total Sorts

- The total number of sorts that ran. at the database or application level, use this value with the Sort Overflows attribute to calculate the percentage of sorts that need more heap space. You can also use it with the Total Sort Time attribute to calculate the average sort time. If the number of sort overflows is small with respect to the total sorts, increasing the sort heap size might have little impact on performance, unless this buffer size is increased substantially. The following value is valid:
- The type is int.
- The unit is sorts.

#### UID SQL Stmts

- The number of SQL UPDATE, INSERT, and DELETE statements that ran. Use this attribute to determine the level of database activity at the application or database level. You can also use the following formula to determine the ratio of UPDATE, INSERT, and DELETE statements to the total number of statements:
  1. Add the number of static SQL statements attempted and the dynamic SQL statements attempted.
  2. Divide the number of UPDATE/INSERT/DELETE SQL statements that ran by the sum derived in step 1.

The following value is valid:

- The type is int.
- The unit is statements.

#### Pool Async Read Time

- The total elapsed time spent reading by database manager prefetchers. Use this attribute to calculate the elapsed time for synchronous reading, using the following formula: total buffer pool physical read time - buffer pool synchronous read time You can also use this attribute to calculate the average asynchronous read time using the following formula: buffer pool asynchronous read

time / buffer pool asynchronous data reads These calculations can be used to understand the I/O work being performed.

- The type is int.
- The unit is milliseconds.

#### Pool Async Data Read Reqs

- The number of asynchronous read requests. To calculate the average number of data pages read per asynchronous request, use the following formula: buffer pool asynchronous data reads / buffer pool asynchronous read requests This average can help to determine the amount of asynchronous I/O in each interaction with the prefetcher.
- The type is int.
- The unit is reads.

### **Component: Buffer Pool**

Information about buffer pool activities.

#### **Dimensions**

##### Buffer Pool Input DB Alias

- The alias of the database provided when calling the snapshot function. The value format is a simple text string with a maximum of 60 characters. Use this attribute to help you identify the specific database to which the monitor data applies. It contains blanks unless you requested monitor information related to a specific database.
- The type is string.

##### Buffer Pool BP Name

- The name of the buffer pool. A new database has a default buffer pool (named IBMDEFAULTBP). The size of the default buffer pool is determined by the platform. Depending on your needs you might choose to create several buffer pools, each of a different size, for a single database. The CREATE, ALTER, and DROP BUFFERPOOL statements allow you to create, change, or remove a buffer pool.
- The type is string.

##### Buffer Pool DB Path

- The full path of the location where the database is stored on the monitored system. Use this attribute with the Database Name attribute to identify the specific database to which the data applies.
- The type is string.

##### Buffer Pool DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If a db partition filter is not specified, data is returned for the current database partition. If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range, the following values are also valid:
- The type is string. This is a key dimension.

##### Buffer Pool Instance Name

- The name of the monitored DB2 instance.
- The type is string. This is a key dimension.

#### Buffer Pool Node Name

- The format is instanceid:hostname:UD for all operating systems. The format for version 6, release 1 of the DB2 agent on Windows systems is instanceid:hostname:UD; on UNIX and Linux systems, the format is instanceid:hostname.
- The type is string.

#### Buffer Pool DB Name

- The real name of the database for which information is collected or to which the application is connected. This name was given to the database when it was created. The value format is a simple text string with a maximum of 60 characters. Use this attribute to identify the specific database to which the data applies.
- The type is string. This is a key dimension.

### Metrics

#### Buffer Pool Pool Sync Data Writes

- The total number of physical write requests that were performed synchronously (that is, physical data page writes that were performed by database manager agents). This value is derived by subtracting the value of the Pool Async Data Writes attribute from the value of the Pool Data Writes attribute. By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing.
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Data L Reads

- The number of logical read requests for data pages that have gone through the buffer pool. This count includes accesses to the following data:
  - Data that is already in the buffer pool when the database manager needs to process the page.
  - Data that is read into the buffer pool before the database manager can process the page.

By using the Pool Data Physical Reads attribute, you can calculate the data page hit ratio for the buffer pool as follows:  $1 - (\text{buffer pool data physical reads} / \text{buffer pool data logical reads})$  By using the Pool Data Physical Reads, Pool Index Physical Reads, and Pool Index Logical Reads attributes, you can calculate the overall buffer pool hit ratio as follows:  $1 - ((\text{buffer pool data physical reads} + \text{buffer pool index physical reads}) / (\text{buffer pool data logical reads} + \text{buffer pool index logical reads}))$  Increasing buffer pool size generally improves the hit ratio until you reach a point of diminishing returns.

- The type is int.
- The unit is reads.

#### Buffer Pool Avg Sync Write Time

- The average elapsed time used to perform a synchronous write. This value is derived by dividing the value of the Pool Sync Write Time attribute by the value of the Pool Sync Write attribute. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device.
- The type is int.
- The unit is milliseconds.

#### Buffer Pool Logical Read Per Min

- The number of logical read operations that are performed on the buffer pool per minute.
- The type is int.



- The unit is reads/minute.

#### Buffer Pool Direct Reads

- The number of read operations that do not use the buffer pool. Use the following formula to calculate the average number of sectors that are read by a direct read:  $\text{direct reads from database} / \text{direct read requests}$ . When using system monitors to track I/O, this data attribute helps to distinguish database I/O from non-database I/O on the device.
- The type is int.
- The unit is reads.

#### Buffer Pool Pool Index Writes

- The number of times a buffer pool index page was physically written to disk. If a buffer pool index page is written to disk for a high percentage of Buffer Pool Index Physical Reads, performance might improve by increasing the number of buffer pool pages available for the database. If all applications are updating the database, increasing the size of the buffer pool might have minimal impact on performance; most pages contain updated data that must be written to disk.
- The type is int.
- The unit is writes.

#### Buffer Pool Pool Async Write Time

- The total elapsed time spent writing data or index pages from the buffer pool to disk by database manager page cleaners. Calculate the elapsed time spent writing pages synchronously by subtracting the value of the Pool Async Write Time attribute from the value of the Pool Physical Write Time attribute. You can also use this attribute to calculate the average asynchronous read time:
  1. Sum the value of the Pool Async Data Writes attribute and the value of the Pool Async Index Writes attribute.
  2. Divide the value of the Pool Async Write Time attribute by the sum from step 1.

These calculations can be used to understand the I/O work being performed.

- The type is int.
- The unit is milliseconds.

#### Buffer Pool Pool Hit Ratio

- The buffer pool hit ratio (as a percentage). The sum of the Pool Data Logical Reads and Pool Index Logical Reads attributes is divided by the value of the Pool Total Reads attribute to derive the pool hit ratio. This attribute can determine whether buffer pool assignment is efficient. If the pool hit ratio is low, increasing the number of buffer pool pages might improve performance.
- The type is double.
- The unit is percent.

#### Buffer Pool Pool Sync Write

- The total number of synchronous index writes. The value is derived by adding the values of the Pool Sync Data Writes attribute and Pool Sync Index Writes attribute.
- The type is int.
- The unit is writes.

#### Buffer Pool Avg Pool Read Time

- The average elapsed time for a read request. This value is derived by dividing the value of the Pool Read Time attribute by the value of the Pool Total Reads attribute. This average is

important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device.

- The type is int.
- The unit is ?.

#### Buffer Pool Pool Index from Estore

- Number of buffer pool index pages copied from extended storage. Required index pages are copied from extended storage to the buffer pool. The copy process might incur the cost of connecting to the shared memory segment, but it saves the cost of a disk read.
- The type is int.
- The unit is pages.

#### Buffer Pool Avg Direct Write Time

- The average elapsed time for a direct write request. This value is calculated by dividing the value of the Direct Write Time attribute by the value of the Direct Writes attribute. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device. The following value is also valid:
- The type is int.
- The unit is result.

#### Buffer Pool Async Write Ratio

- The ratio of buffer pool asynchronous data writes to the total number of pool writes for the database.
- The type is int.
- The unit is ratio.

#### Buffer Pool Direct Write Reqs

- The number of requests to perform a direct write of one or more sectors of data. Use the following formula to calculate the average number of sectors that are written by a direct write: direct writes to database / direct write requests The following value is also valid:
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Index L Reads

- The number of logical read requests for index pages that have gone through the buffer pool. This count includes accesses to the following index pages:
  - Pages that are already in the buffer pool when the database manager needs to process the page.
  - Pages that are read into the buffer pool before the database manager can process the page.

By using the Buffer Pool Index Physical Reads attribute, you can calculate the index page hit ratio for the buffer pool as follows:  $1 - (\text{buffer pool index physical reads} / \text{buffer pool index logical reads})$  If the hit ratio is low, increasing the number of buffer pool pages might improve performance.

- The type is int.
- The unit is requests.

#### Buffer Pool Avg Data Page Read per Async Req

- The average number of pages read for each asynchronous request. This value is derived by dividing the value of the Pool Async Data Reads attribute by the value of the Pool Async Data Read Reqs attribute.
- The type is int.
- The unit is pages.

#### Buffer Pool Pool Sync Index Reads

- The number of index pages read synchronously (that is, physical index page reads that were performed by database manager agents) into the buffer pool. This value is derived by subtracting the value of the Pool Async Index Reads attribute from Pool Index Physical Reads attribute. By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working.
- The type is int.
- The unit is pages.

#### Buffer Pool Pool Data to Estore

- Number of buffer pool data pages copied to extended storage. Pages are copied from the buffer pool to extended storage when they are selected as victim pages. As a result of the copying process, there is sufficient space for new pages in the buffer pool.
- The type is int.
- The unit is pages.

#### Buffer Pool Avg Sync Read Time

- The average elapsed time used to perform a synchronous read. This value is derived by dividing the value of the Pool Sync Read Time attribute by the value of the Pool Sync Read attribute. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device.
- The type is int.
- The unit is milliseconds.

#### Buffer Pool Direct Write Time

- The elapsed time (in milliseconds) required to perform the direct writes. Use the following formula to calculate the average direct write time per sector:  $\text{direct write time} / \text{direct writes to database A}$  high average time might indicate an I/O conflict.
- The type is int.
- The unit is milliseconds.

#### Buffer Pool Pool Sync Read Time

- The elapsed time used to perform all synchronous reads. This value is derived by subtracting the value of the Pool Async Read Time attribute from the value of the Pool Read Time attribute. Use this attribute to understand the I/O work being performed.
- The type is int.
- The unit is milliseconds.

#### Buffer Pool Pool Sync Write Time

- The total elapsed time used to perform all synchronous writes. This value is derived by subtracting the value of the Pool Async Write Time attribute from the value of the Pool Write Time attribute.
- The type is int.
- The unit is milliseconds.

#### Buffer Pool Files Closed

- The total number of database files closed. The database manager opens files for reading and writing into and out of the buffer pool. The maximum number of database files open by an application at any time is controlled by the MAXFILOP configuration parameter. If the maximum is reached, one file is closed before the new file is opened. Note that the actual number of files opened might not equal the number of files closed. The following value is also valid:
- The type is int.
- The unit is files.

#### Buffer Pool Pool Data Writes

- The number of times a buffer pool data page was physically written to disk. A buffer pool data page is written to disk for the following reasons:
  - To free a page in the buffer pool so another page can be read
  - To flush the buffer pool.

If a buffer pool data page is written to disk for a high percentage of Buffer Pool Data Physical Reads, performance might improve by increasing the number of buffer pool pages available for the database.

- The type is int.
- The unit is writes.

#### Buffer Pool Pool Data from Estore

- Number of buffer pool data pages copied from extended storage. Required pages are copied from extended storage to the buffer pool. The copy process might incur the cost of connecting to the shared memory segment, but it saves the cost of a disk read. The following value is also valid:
- The type is int.
- The unit is pages.

#### Buffer Pool Pool Async Data Read Reqs

- The number of asynchronous read requests. To calculate the average number of data pages read per asynchronous request, use the following formula:  $\text{buffer pool asynchronous data reads} / \text{buffer pool asynchronous read requests}$  This average can help to determine the amount of asynchronous I/O in each interaction with the prefetcher.
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Async Read Time

- The total elapsed time spent reading by database manager prefetchers. Use this attribute to calculate the elapsed time for synchronous reading, using the following formula:  $\text{total buffer pool physical read time} - \text{buffer pool synchronous read time}$  You can also use this attribute to calculate the average asynchronous read time using the following formula:  $\text{buffer pool asynchronous read time} / \text{buffer pool asynchronous data reads}$  These calculations can be used to understand the I/O work being performed.
- The type is int.
- The unit is milliseconds.

#### Buffer Pool Pool Async Data Reads

- The number of pages read asynchronously into the buffer pool. Use this attribute with the Buffer Pool Data Physical Reads attribute to calculate the number of physical reads that were

performed synchronously (that is, physical data page reads that were performed by database manager agents). Use the following formula:  $\text{buffer pool data physical reads} - \text{buffer pool synchronous data reads}$ . By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working.

- The type is int.
- The unit is pages.

#### Buffer Pool Pool Sync Index Writes

- The number of physical index write requests that were performed synchronously (that is, physical index page writes that were performed by database manager agents). This value is derived by subtracting the value of the Pool Async Index Writes attribute from the value of the Pool Index Writes attribute. By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing.
- The type is int.
- The unit is requests.

#### Buffer Pool Direct Read Time

- The elapsed time (in milliseconds) required to perform the direct reads. Use the following formula to calculate the average direct read time per sector:  $\text{direct read time} / \text{direct reads}$  from database A. A high average time might indicate an I/O conflict.
- The type is int.
- The unit is milliseconds.

#### Buffer Pool Avg Pool Write Time

- The average elapsed time for a write request. This value is derived by dividing the value of the Pool Write Time attribute by the value of the Pool Total Writes attribute.
- The type is int.
- The unit is milliseconds.

#### Buffer Pool Prefetch Ratio

- The percentage of asynchronous read operations that the prefetcher performed for sequential scans.
- The type is int.
- The unit is percent.

#### Buffer Pool Pool Index P Reads

- The number of physical read requests to get index pages into the buffer pool. see the Pool Index Logical Reads attribute for information about how to use this element.
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Sync Read

- The total number of synchronous reads. This value is derived by adding the values of the Pool Sync Data Reads and Pool Sync Index Reads attributes.
- The type is int.
- The unit is reads.

#### Buffer Pool Pool Index to Estore

- Number of buffer pool index pages copied to extended storage. Pages are copied from the buffer pool to extended storage when they are selected as victim pages. As a result of the copying process, there is sufficient space for new pages in the buffer pool.
- The type is int.
- The unit is pages.

#### Buffer Pool Pool Sync Data Reads

- The number of physical data page reads that were performed by database manager agents. This value is derived by subtracting the value of the Pool Async Data Reads attribute from the Pool Data Physical Reads attribute. By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working.
- The type is int.
- The unit is reads.

#### Buffer Pool Pool Async Index Reads

- The number of index pages read asynchronously into the buffer pool by a prefetcher. Asynchronous reads are performed by database manager prefetchers. Use this attribute with the Buffer Pool Index Physical Reads attribute to calculate the number of physical reads that were performed synchronously (that is, physical index page reads that were performed by database manager agents). Use the following formula: buffer pool index physical reads - buffer pool asynchronous index reads. By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working.
- The type is int.
- The unit is reads.

#### Buffer Pool Pool Total Writes

- The total number of write requests. This attribute is the total of the Pool Data Writes and Pool Index Writes attributes. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal.
- The type is int.
- The unit is writes.

#### Buffer Pool Avg Direct Read Time

- The average elapsed time for a direct read request. This value is calculated by dividing the value of the Direct Read Time attribute by the value of the the Direct Reads attribute. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device.
- The type is int.
- The unit is ?.

#### Buffer Pool Pool Total Reads

- The total number of read requests that required I/O to get data pages and index pages into the buffer pool. This attribute is the total of the Pool Data Physical Reads and Pool Index Physical Reads attributes. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal. The following value is valid:
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Async Data Writes

- The number of times a buffer pool data page was physically written to disk by an asynchronous page cleaner or by a prefetcher. A prefetcher might have written dirty pages to disk to make

space for the pages being prefetched. Use this attribute with the Buffer Pool Data Writes attribute to calculate the number of physical write requests that were performed synchronously (that is, physical data page writes that were performed by database manager agents). Use the following formula:  $\text{buffer pool data writes} - \text{buffer pool asynchronous data writes}$  By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing.

- The type is int.
- The unit is writes.

#### Buffer Pool Pool Data P Reads

- The number of read requests that required I/O to get data pages into the buffer pool.
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Async Index Writes

- The number of times a buffer pool index page was physically written to disk by an asynchronous page cleaner or a prefetcher. A prefetcher might have written dirty pages to disk to make space for the pages being prefetched. Use this attribute with the Buffer Pool Index Writes attribute to calculate the number of physical index write requests that were performed synchronously. That is, physical index page writes that were performed by database manager agents. Use the following formula:  $\text{buffer pool index writes} - \text{buffer pool asynchronous index writes}$  By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing.
- The type is int.
- The unit is writes.

#### Buffer Pool Direct Writes

- The number of write operations that do not use the buffer pool. Use the following formula to calculate the average number of sectors that are written by a direct write:  $\text{direct writes to database} / \text{direct write requests}$  When using system monitors to track I/O, this data attribute helps to distinguish database I/O from non-database I/O on the device.
- The type is int.
- The unit is writes.

#### Buffer Pool Pool Write Time

- The total amount of time spent physically writing data or index pages from the buffer pool to disk. Use this attribute with the Buffer Pool Data Writes and Buffer Pool Index Writes attributes to calculate the average page-write time. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device. The following value is valid:
- The type is int.
- The unit is milliseconds.

#### Buffer Pool Application Direct Read Reqs

- The number of requests to perform a direct read of one or more sectors of data. Use the following formula to calculate the average number of sectors that are read by a direct read:  $\text{direct reads from database} / \text{direct read requests}$  The following value is also valid:
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Read Time

- The total amount of elapsed time spent processing read requests that caused data or index pages to be physically read from buffer pool to disk. Use this attribute with the Buffer Pool Data Physical Reads and Buffer Pool Index Physical Reads attributes to calculate the average page-read time. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device. The following value is also valid:
- The type is int.
- The unit is milliseconds.

### **Component: Database Efficiency**

Information about the efficiency of the database and identify any problem areas for corrective action. All values are integers that are calculated from the first application connection, unless otherwise noted.

#### **Dimensions**

##### Efficiency Catalog Cache Size

- The value in units of 4-KB pages of the catalog cache size. This value is the maximum amount of space that the catalog cache can use from the database heap (dbheap). The catalog cache is referenced whenever a table, view, or alias name is processed during the compilation of an SQL statement. It is dynamically allocated from dbheap, as required, until the catalog cache size is reached.
- The type is int.

##### Efficiency System Tablespaces

- The number of SMS tablespaces in the database. Use the returned value to evaluate the use of SMS tablespaces and their effects on performance. Table data that is read from disk is available in the database buffer pool. Sometimes a data page is freed from the buffer pool before it is used. For SMS tablespaces, when the database manager requests that data page from the file system, the data page might still be in the cache of the file system. Having the page in the cache saves an input and output operation that would otherwise have been required. (For more information, see the DB2 administration documentation for the version of DB2 that you are using.) If you have many SMS tablespaces, you can increase the size of the file system cache to take advantage of this extra buffering.
- The type is int.

##### Efficiency Buff Page

- Use the returned value (in units of 4-KB pages) to analyze the input and output work being performed for the database. Synchronous input and output operations for a database are performed by database manager agents. Asynchronous input and output operations are performed by prefetchers (reads) and page cleaners (writes). In general, asynchronous input and output helps your applications run faster. In the currently supported releases of DB2, multiple buffer pools might be defined in a single database. For instance, buffer pools can be defined and associated with a particular tablespace. Each buffer pool created can be given its own individual size. The buffpage attribute serves only as a default value for buffer pools created within a particular database. Therefore, the value of the buffpage attribute is much less critical to performance in current releases of DB2, because most buffer pools are given an individual size when created. The buffpage attribute must not be used to evaluate or tune the performance of DB2 unless it is used as the default value when creating buffer pools in a database.
- The type is int.

##### Efficiency Max Active Applications



- The value of the maximum number of active applications. This value is the maximum number of concurrent applications that can be connected (both local and remote) to a database. Because each application that attaches to a database causes some private memory to be allocated, allowing a large number of concurrent applications potentially uses more memory. Increasing the value of this parameter without lowering the maxlocks parameter or increasing the locklist parameter can cause you to reach the database limit on locks (locklist) rather than the application limit. The result can be pervasive lock escalation problems.
- The type is int.

#### Efficiency Appl Heap Size

- The size (in 4-KB pages) of the application heap that is available for each individual agent in the database during the monitoring interval. Increase the value of the parameter if your application receives an error indicating that there is not enough storage in the application heap. The heap is allocated when an agent or subagent is initialized for an application. The amount allocated is the minimum amount needed to process the request given to the agent or subagent. When the agent or subagent requires more heap space to process larger SQL statements, the database manager allocates memory as needed, up to the maximum specified by the parameter.
- The type is int.

#### Efficiency Change Pages Threshold

- The value in percentage units of the changed pages threshold. This value sets a limit on how much buffer pool space can be occupied by changed pages before the asynchronous page cleaners are started, if they are not currently active. Asynchronous page cleaners write changed pages from the buffer pool to disk before the space in the buffer pool is required by a database agent. This means that the agents do not need to wait for a changed page to be written out before being able to read a page, and application transactions run faster.
- The type is int.

#### Efficiency Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string.

#### Efficiency Min Commit

- The value of the number of commits to group. By using this parameter you can delay the writing of log records to disk until a minimum number of commits have been performed. This delay can help reduce the overhead associated with writing log records and can improve performance. The default value for mincommit is 1, which can be too low for your environment. By sampling the number of transactions per-second throughout the day, you can determine the peak per second rate and adjust the value of the mincommit parameter to accommodate all or most transactions. This adjustment minimizes the number of log writes under the heaviest conditions. As you increase the value of the mincommit parameter, you might also need to increase the log buffer size (LOGBUFSZ parameter) to avoid filling the log buffer. Filling the log buffer also forces the writing of log records to disk. If you change the value of the mincommit parameter, you must change the value for the logbufsz configuration parameter.
- The type is int.

#### Efficiency Tablespaces

- The number of tablespaces in the database. Use this attribute to track database growth over a period of time.
- The type is int.

#### Efficiency Tablespaces Long Data

- The number of tablespaces that store LONG data in the database. Use this attribute to track database growth over a period of time. LONG data can take up a large amount of space in a database.
- The type is int.

#### Efficiency Sequential Detect

- The current value of the sequential detection flag, which determines if the database manager must perform sequential detection. The database manager can monitor input and output operations. If sequential page reading is occurring, the database manager can activate input and output prefetching. This type of sequential prefetch is known as sequential detection. If this configuration parameter is set to no , prefetching takes place only if the database manager determines that it is useful (for example, in table sorts).
- The type is int.

#### Efficiency Sort Heap

- The current value in units of 4-KB pages of the sort heap size. This value is the maximum amount of memory that can be allocated as sort heap for each sort within a database. The sort heap is the memory block where data is sorted. The following value is valid:
- The type is int.

#### Efficiency Lock List

- The value in units of 4-KB pages of the maximum storage for lock lists. This value is the amount of storage that is allocated to the lock list. There is one lock list for each database, and it contains the locks held by all applications concurrently connected to the database. Too small a value can lead to excessive lock waits. Too high a value compared to normal operating levels can deprive the system of resources or memory.
- The type is int.

#### Efficiency DB Tablespaces

- The number of Database Managed Space tablespaces in the database. Use this attribute to track database growth over a period of time. The following value is also valid:
- The type is int.

#### Efficiency Log Buffer Size

- This value specifies the amount of the database heap to use as a buffer for log records before writing these records to disk. It is important that the log buffer can hold the amount of log space used by an average transaction. Otherwise, logging performance decreases and slows the overall system.
- The type is int.

#### Efficiency Snapshot Timestamp

- The date and time when the database system monitored information was collected. Use this attribute to help correlate data chronologically if you are saving the results in a file or database for ongoing analysis.
- The type is timestamp.

#### Efficiency Max Locks

- The value of the maximum percentage of lock list before escalation. This value specifies the percentage of the lock list that an application can hold before the database manager performs lock escalation. Lock escalation can increase contention, which reduces system throughput and increases user response time. The values for the maxlocks and maxappls parameters must

satisfy (MAXLOKS x MAXAPPLS) >100, and each lock uses 32 bytes. Rebind application packages after changing this parameter.

- The type is int.

#### Efficiency Num IO Servers

- The current value of the number of input and output servers. This value specifies the number of input and output servers for a database. Input and output servers are used on behalf of the database agents to perform asynchronous input and output operations for utilities such as backup and restore, and to perform prefetch input and output (in which case, they are called prefetchers) operations. Prefetchers read pages from disk into the buffer pool in anticipation of their use. In most situations, these pages are read just before they are needed. However, prefetchers can cause unnecessary input and output operations by reading pages into the buffer pool that might not be used. For example, an application starts reading through a table, and prefetchers read consecutive pages into the buffer pool before the pages are required by the application. Then the application fills the application buffer and stops reading. Meanwhile, the prefetchers already have performed the input and output operations for additional pages and the buffer pool is partially taken up with those pages. To exploit all the input and output devices in the system, a good value for num\_ioservers to use is one or two more than the number of physical devices on which the database is established.
- The type is int.

#### Efficiency Instance Name

- The name of the monitored DB2 instance.
- The type is string. This is a key dimension.

#### Efficiency Database Efficiency Log Primary

- The number of primary log files.
- The type is int.

#### Efficiency Database Efficiency New Log Path

- The current value of the newlogpath configuration parameter. You use the newlogpath configuration parameter to specify a new location for the log files. The specified path does not become the current log path until both of the following conditions are met:
  - The database is in a consistent state.
  - All users are disconnected from the database.

When the first new connection is made to the database, the database manager moves the logs to this location.

- The type is string.

#### Efficiency DB Name

- The real name of the database for which information is collected or to which the application is connected. This name was given to the database when it was created. The value format is a simple text string with a maximum of 60 characters.
- The type is string. This is a key dimension.

#### Efficiency DB Capture Lag

- The time difference in minutes between the current timestamp and the last timestamp recorded by the Capture program. This time difference is the Capture lag. Use the returned value to determine whether the Capture program is keeping up with the DB2 database log. The Capture program uses an interface to the DB2 database log or journal to detect and save changes to the data in the tables registered for replication.

- The type is int.

#### Efficiency Appl Control Heap Size

- The maximum size (in 4-KB pages) for the application control heap in the database during the monitoring interval. The heap is required to share information among agents working on behalf of the same application at a node in a massively parallel processing (MPP) or a symmetric multiprocessor (SMP) system. If complex applications are being run or the MPP configuration has a large number of nodes, you must increase the size of this heap. In a partitioned database environment, this heap is used to store copies of the executing section of SQL statements for agents and subagents. However, symmetric multiprocessor agents (SMP), subagents, and agents in all other environments use appl heap size.
- The type is int.

#### Efficiency Package Cache Size

- The current value in units of 4-KB pages of the package cache size. This value controls the amount of application heap memory to be used for caching static and dynamic SQL statements of a package. You must experiment with the size of the package cache to find the optimal number for this attribute. For example, you can use a smaller package cache size if there is no increase in the number of package cache inserts when you decrease the size of the cache. Decreasing the package cache size frees up system resources for other work. However, increasing the package cache size can improve overall system performance if it results in a decrease of package cache inserts.
- The type is int.

#### Efficiency Database Heap

- The value in units of 4-KB pages of the database heap. This value is the maximum amount of memory allowed for a database heap. There is one database heap for each database. It is used on behalf of all applications connected to the database. Refining dbheap has minimal impact on performance. The main function of this parameter is to prevent the database manager from allocating an excessive amount of space for a particular database.
- The type is int.

#### Efficiency Tables

- The number of tables in the database. Use this attribute to track database growth due to an increased number of tables over a period of time.
- The type is int.

#### Efficiency Number of I/O Cleaners

- The current value of the number of asynchronous page cleaners. This parameter specifies the number of asynchronous page cleaners for a database. Page cleaners monitor the buffer pool and asynchronously write out changed pages to disk to free space in the buffer pool.
- The type is int.

#### Efficiency Database Efficiency Restore Pending

- The RESTORE PENDING status in the database during the last monitoring interval.
- The type is int.

#### Efficiency DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If a db partition filter is not specified, data is returned for the current database partition. If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data collection includes both

aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range, the following values are also valid:

- The type is string. This is a key dimension.

## Metrics

### Efficiency Database Efficiency Pool I/O per Sec

- The rate (per second) of buffer pool input and output operations for the database. Buffer pool input and output includes all physical data and index pages that go through the buffer pool when read or written. Use the returned value to determine how efficient your data storage device is. A low value indicates the presence of an input and output wait, in which case you must move data to a different device. The following value is valid:
- The type is int.
- The unit is operations/second.

### Efficiency Estore Read/Write Ratio for Interval

- The ratio as a percentage of data and index pages copied from extended storage to pages copied to extended storage during the monitoring interval. When a page is transferred from extended storage to the buffer pool, you save a system input and output call. However, you still incur the cost of attaching to the extended memory segment, copying the page, and detaching from the segment. Use the returned value to determine if you would benefit from using extended storage. The higher the ratio, the more likely you are to benefit. In general, extended storage is particularly useful if input and output activity is very high on your system. The following value is also valid:
- The type is double.
- The unit is reads/writes.

### Efficiency Lock Waits Percent

- The percentage of currently connected applications that are waiting for a lock in the database. The value is derived through this formula:  $100 * \text{locks waiting} / \text{appls cur cons}$  If the returned value is high compared to normal operating levels, the applications can have concurrency problems. You must identify applications that are holding locks or exclusive locks for long periods of time and determine whether they can release their locks more often.
- The type is double.
- The unit is percent.

### Efficiency SQL Stmts Rate for Interval

- The rate, in issued SQL statements per second, at which SQL statements that run during the monitoring interval. The value format is an integer.
- The type is double.
- The unit is percent.

### Efficiency Database Efficiency Pool Sync Data Writes

- The total number of physical write requests that were performed synchronously (that is, physical data page writes that were performed by database manager agents). This value is derived by subtracting the value of the Pool Async Data Writes attribute from the value of the Pool Data Writes attribute. By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing.
- The type is int.
- The unit is writes.

### Efficiency Avg Pool Async Data Reads

- The average number of buffer pool asynchronous data reads when compared to the total number of pool reads for the database. The value is derived through this formula:  $\text{pool async data reads} / (\text{pool data p reads} + \text{pool index p reads})$ . Use the returned value to gain insight into how well the prefetchers are working and to refine the `num_ioservers` configuration parameter. If the returned value is low compared to normal operating levels, there might not be enough input and output servers to prefetch data into the buffer, causing the database manager agents to spend extra time on physical reads. Increase the number of input and output servers by increasing the value of the `num_ioservers` configuration parameter. If too many servers are allocated, system performance is not reduced because the extra input and output servers are not used.
- The type is int.
- The unit is reads.

#### Efficiency Database Efficiency Pool Sync Write Time

- The total elapsed time used to perform all synchronous writes. This value is derived by subtracting the value of the Pool Async Write Time attribute from the value of the Pool Write Time attribute.
- The type is int.
- The unit is milliseconds.

#### Efficiency Avg Lock Escal per Conn for Interval

- The average lock escalations per connection for this database during the monitoring interval. The value format is an integer. A lock is escalated when the total number of locks that an application holds reaches the maximum amount of lock list space available to the application, or the lock list space consumed by all applications is approaching the total lock list space. When an application reaches the maximum permitted number of locks and no additional locks can be escalated, the application uses space in the lock list that is allocated for other applications. When the entire lock list is full, an error occurs.
- The type is int.
- The unit is escalations.

#### Efficiency Avg Locks Held

- The average number of locks held by each currently connected application in the database. The value is derived through this formula:  $\text{locks held} / \text{appls cur cons}$ . If the returned value is high compared to normal operating levels, it can indicate that one or more applications is using an excessive number of locks. Refine such applications to improve performance.
- The type is int.
- The unit is locks.

#### Efficiency Pool Hit Ratio Percent for Interval

- The overall buffer pool hit ratio as a percentage for the database during the monitoring interval. This hit ratio includes both index and data page activity. The overall buffer pool hit ratio indicates the percentage of page requests for which the database manager did not need to load a page from disk to service. (That is, the page was already in the buffer pool.) The greater the buffer pool hit ratio, the lower the frequency of disk input and output. If the hit ratio is low compared to normal operating levels, increasing the number of buffer pool pages can improve performance. A ratio of zero indicates that pages needed to be read for every request. For a large database, increasing the buffer pool size can have a minimal effect on the buffer pool hit ratio. Such a database can have so large a number of data pages that the statistical chance of a hit is not increased by an increase of the buffer pools. However, even though the data might be too large to fit in the buffer pool, the entire index can fit. In this case, you can refine buffer pool sizes until the overall buffer pool hit ratio stops increasing, and then refine the buffer pool until the buffer pool index hit ratio no longer increases.

- The type is double.
- The unit is percent.

#### Efficiency Database Efficiency Primary Log Used Percent

- The percentage of total log space used by the primary log. Use the returned value to help you evaluate the allocated amount of primary log space and refine the log buffer size, log file size, and primary log configuration parameters. The returned value is valid only if circular logging is used.
- The type is double.
- The unit is percent.

#### Efficiency Triggers

- The number of triggers defined in the database. Use this attribute to track the use of triggers in the database. There are benefits to using triggers, including faster application development, easier maintenance, and global enforcement of business rules. For more information, see the DB2 administration documentation for the version of DB2 that you are using.
- The type is int.
- The unit is triggers.

#### Efficiency Database Efficiency Pool Sync Read Time

- The elapsed time used to perform all synchronous reads. This value is derived by subtracting the value of the Pool Async Read Time attribute from the value of the Pool Read Time attribute. Use this attribute to understand the I/O work being performed.
- The type is int.
- The unit is milliseconds.

#### Efficiency Primary Log Used Top

- The maximum bytes of primary logs used.
- The type is int.
- The unit is bytes.

#### Efficiency Database Efficiency Avg Direct Write Time

- The average time in milliseconds for performing direct writes to the database. A high average time can indicate the existence of an input and output conflict. The value is derived through this formula:  $\text{direct write time} / \text{direct write}$ . The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Efficiency Database Efficiency Pool Sync Write

- The total number of synchronous index writes. The value is derived by adding the values of the Pool Sync Data Writes attribute and Pool Sync Index Writes attribute.
- The type is int.
- The unit is writes.

#### Efficiency Database Efficiency Current Primary Log Used Percent

- The percentage of primary log space that is currently in use.
- The type is double.
- The unit is percent.

#### Efficiency Database Efficiency Avg Direct Read Time

- The average time in milliseconds that is used to perform direct reads to the database. The value is derived through this formula:  $\text{direct read time} / \text{direct reads}$ . The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Efficiency Internal Auto Rebinds

- The number of automatic rebinds or recompiles that were attempted in the database. Use the returned value to determine the level of database activity. Automatic rebinds are the internal binds that the system performs when a package is invalidated. They can have a significant impact on performance and must be minimized where possible.
- The type is int.
- The unit is rebinds.

#### Efficiency User Indexes

- The number of indexes created by users in the database. Indexes created by SYSIBM are not counted. Use this to track the use of indexes in the database. The use of indexes can improve performance; for example, faster sorting of data. However, indexes can also have adverse effects on performance; for example, each INSERT or DELETE operation performed on a table requires additional updating of each index on that table. For a discussion of this topic, see the DB2 administration documentation for the version of DB2 that you are using.
- The type is int.
- The unit is indexes.

#### Efficiency Total Direct I/O Time

- The total time in milliseconds applied to direct reads and writes for the database. The returned value indicates the amount of time that the database performs direct reads and writes. A high returned value compared to normal operating levels can indicate the presence of an input and output conflict.
- The type is int.
- The unit is milliseconds.

#### Efficiency Avg Pool Writes per Read

- The ratio of total pool writes to pool reads for the database. The value is derived through this formula:  $(\text{pool data writes} + \text{pool index writes}) / (\text{pool data p reads} + \text{pool index p reads})$ . If the returned value is greater than 1, you can improve performance by increasing the available buffer pool space. A returned value greater than 1 indicates that at least one write to disk had to occur (either to free a page in the buffer pool, or to flush the buffer pool) before a page can be read into the buffer pool. You can increase the available buffer pool space by freeing the space more often or by increasing the total space for the buffer pool. The following value is also valid:
- The type is int.
- The unit is writes/reads.

#### Efficiency Internal Deadlock Rollbacks Percent for Interval

- The percentage of rollbacks that were due to deadlock during the monitoring interval. Use the returned value to distinguish those rollbacks caused by internal deadlocks from rollbacks caused by other situations (for example, incomplete imports). The returned value is the percentage of internal rollbacks due to internal deadlocks since the first database connection or the last reset of the database monitor counters.
- The type is double.
- The unit is percent.



#### Efficiency Days Since Last Backup

- The numbers of days since the last database backup was completed. The value format is an integer. The value for no backup completed is 2147483647. The following value is also valid:
- The type is int.
- The unit is days.

#### Efficiency DB Heap Top

- This data attribute (now maintained for DB2 version compatibility) measures memory usage, but not exclusively usage by the database heap.
- The type is int.
- The unit is bytes.

#### Efficiency Database Efficiency Current Secondary Log Used Percent

- The percentage of secondary log space that is currently in use.
- The type is double.
- The unit is percent.

#### Efficiency Database Efficiency Pool Sync Index Reads

- The number of pool index physical reads minus the pool async index reads.
- The type is int.
- The unit is reads.

#### Efficiency Appls in DB2

- The number of applications currently executing in the database. The following value is also valid:
- The type is int.
- The unit is applications.

#### Efficiency Avg Pages per Cleaner for Interval

- The average number of pages written per page cleaner that are invoked for the database during the monitoring interval. Use the returned value to determine how many pages are handled by the page cleaners of this database. If this value increases over time, you can define more page cleaners. The following value is also valid:
- The type is int.
- The unit is pages/cleaner.

#### Efficiency Database Efficiency Avg Sync I/O Time

- The average time (in milliseconds) to perform synchronous input and output operations for the database. Use the returned value to analyze the input and output work being performed for the database. Synchronous input and output operations for a database are performed by database manager agents. Asynchronous input and output operations are performed by prefetchers (reads) and page cleaners (writes). In general, asynchronous input and output helps your applications run faster. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Efficiency Rollback Rate for Interval

- The rate, in rollbacks per second, at which unit-of-work rollbacks were attempted during the monitoring interval. Unit-of-work rollbacks include SQL ROLLBACK statements that are issued from applications and INTERNAL ROLLBACKS that the database manager initiates.
- The type is double.
- The unit is percent.

#### Efficiency Invalid Triggers

- The number of triggers that are marked not valid in the database. Use the returned value to determine the number of triggers that must be revalidated. A trigger is marked not valid if an object on which the trigger depends is dropped. To revalidate such a trigger, retrieve its definition from the database system catalog and submit a new CREATE TRIGGER statement.
- The type is int.
- The unit is triggers.

#### Efficiency Internal Rows Updated

- The number of rows updated in the database as a result of internal activity. Use the returned value to gain insight into internal activity within the database. If this activity is high compared to normal operating levels, you can evaluate your table design to determine if the referential constraints that you defined are necessary.
- The type is int.
- The unit is rows.

#### Efficiency Database Efficiency Lock List in Use Percent

- The percentage of space used in the locklist of this database. Use the returned value to determine how much of the locklist space is free for new locks to be requested.
- The type is double.
- The unit is percent.

#### Efficiency Database Efficiency Avg Sync Read Time

- The average elapsed time used to perform a synchronous read. This value is derived by dividing the value of the Pool Sync Read Time attribute by the value of the Pool Sync Read attribute. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device.
- The type is int.
- The unit is milliseconds.

#### Efficiency Database Efficiency Deadlocks for Interval

- The number of deadlocks detected in the database during the monitoring interval. Use the returned value to determine whether applications are experiencing conflict problems in the database. You can resolve the problem by determining in which applications the deadlocks are occurring. You can then try to modify the applications to better enable them to run concurrently. The following value is also valid:
- The type is int.
- The unit is deadlocks.

#### Efficiency Cur Cons Percent

- The percentage of applications currently connected.
- The type is double.
- The unit is percent.

### Efficiency Invalid Packages

- The number of all packages that are currently marked not valid in the database. Use the returned value as an indication of the current number of packages that are not valid. A package is marked not valid if it depends on an object (for example, a table) and that object is dropped. The number of packages that are not valid can indicate how many automatic rebinds are necessary in the database. Such packages automatically rebound the next time they are accessed, unless a trigger was dropped or the dropped object was not recreated. Use of automatic rebinds can significantly lower performance, and must be minimized if possible. The following value is also valid:
- The type is int.
- The unit is packages.

### Efficiency Database Efficiency Avg Sect Written per Direct Write

- The average number of sectors that are written by a direct write to the database. The value is derived through this formula:  $\text{direct writes} / \text{direct write reqs}$ . Direct writes do not use the buffer pool, which results in poor performance because the data is physically written from disk each time. If you are using system monitors to track input and output for the device, this value helps you distinguish database input and output from non-database input and output. The following value is also valid:
- The type is int.
- The unit is sectors.

### Efficiency Views

- The number of views in the database. Use this attribute to track the use of views in the database. Views can be created to limit access to sensitive data, while allowing more general access to other data. This provides flexibility in the way your programs and end-user queries can look at the table data.
- The type is int.
- The unit is views.

### Efficiency Database Efficiency Total Log Used

- The total log space used in bytes. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal.
- The type is int.
- The unit is bytes.

### Efficiency Database Efficiency Avg Sect Read per Direct Read

- The average number of sectors that are read by a direct read for the database. The value is derived through this formula:  $\text{direct reads} / \text{direct read reqs}$ . Direct reads do not use the buffer pool, and so result in poor performance because the data is physically read from disk each time. If you are using system monitors to track input and output for the device, this value helps you distinguish database input and output from non-database input and output. The following value is also valid:
- The type is int.
- The unit is sectors.

### Efficiency Database Efficiency DDL SQL Percent for Interval

- The percentage of total SQL statements that were SQL DDL statements during the monitoring interval. Due to the high activity in the system catalog tables, try to keep DDL statement activity to a minimum. If the returned value is high compared to normal operating levels, determine the activity causing it to be high and restrict it from being performed. Examples of DDL statements

are CREATE TABLE, CREATE VIEW, ALTER TABLE, and DROP INDEX. You can also use the returned value to refine the package cache hit ratio for this application. DDL statements can also affect the package cache by invalidating sections that are stored there and causing additional system overhead due to section recompilation.

- The type is double.
- The unit is percent.

#### Efficiency Pool Hit Ratio Index Percent for Interval

- The database index page hit ratio (as a percentage) for the buffer pool during the monitoring interval. The index page hit ratio for the buffer pool indicates the percentage of index page requests for which the database manager did not need to load an index page from disk to service. That is, the index page was already in the buffer pool. The higher the returned value, the lower the frequency of disk input and output, and the faster the performance. If the hit ratio is low compared to normal operating levels, increasing the number of buffer pool pages can improve performance.
- The type is double.
- The unit is percent.

#### Efficiency Sort Overflows Percent for Interval

- The percentage of application sorts that overflowed during the monitoring interval. An overflow occurs when a sort has run out of space in the sort heap and requires disk space for temporary storage. If this percentage is high, you might want to adjust the database configuration by increasing the value of the SORTHEAP configuration parameter. The value format is an integer.
- The type is double.
- The unit is percent.

#### Efficiency lock Timeouts for Interval

- The number of times that a request to lock an object were timed out instead of being granted during the monitoring interval. The value format is an integer.
- The type is int.
- The unit is timeouts.

#### Efficiency Database Efficiency Pool Sync Read

- The total number of synchronous reads. This value is derived by adding the values of the Pool Sync Data Reads and Pool Sync Index Reads attributes.
- The type is int.
- The unit is reads.

#### Efficiency Commit Stmt per Sec

- The total number of commits initiated internally by the database per second. Use the returned value to determine rates of database activity.
- The type is int.
- The unit is commits/second.

#### Efficiency DB Capture Error

- The number of errors encountered by the Capture program within the last five minutes. Use the returned value to determine whether the Capture program encountered an error that prevented it from running. If any errors are detected, the Capture program came down at the time the error occurred. The Capture program might or might not still be down. The Capture program is the most critical replication component in the replication system. If the Capture program is not active, there are no new change records to apply to the target systems. If your data concurrency

requirements are high and you want to ensure that the Capture program runs continuously, use this monitor to determine when the Capture program encounters an error that prevents it from running.

- The type is int.
- The unit is errors.

#### Efficiency Database Efficiency Avg Data Page Read per Async Req

- The average number of pages read for each asynchronous request. This value is derived by dividing the value of the Pool Async Data Reads attribute by the value of the Pool Async Data Read Reqs attribute. Use this attribute to determine whether good enough data pages were read per asynchronous request. The following value is also valid:
- The type is int.
- The unit is pages.

#### Efficiency Internal Deadlock Rollbacks Percent

- The percentage of the total number of internal rollbacks due to deadlocks. Use the returned value to distinguish those rollbacks caused by internal deadlocks from rollbacks caused by other situations (for example, incomplete imports). The returned value is the percentage of internal rollbacks due to internal deadlocks since the first database connection or the last reset of the database monitor counters.
- The type is int.
- The unit is percent.

#### Efficiency Total Pool Phys Read

- The total time in milliseconds applied to processing read requests that caused data or index pages to be physically read from disk to the buffer pool for the database. The value is derived through this formula: pool data p reads + pool index p reads . The returned value is used to calculate the average pool read time. This average can indicate the presence of an input and output wait, which in turn can indicate that you must move data to a different device.
- The type is int.
- The unit is milliseconds.

#### Efficiency Database Efficiency Internal Commits

- The total number of commits initiated internally by the database. Use the returned value to gain insight into internal activity within the database. The returned value is also used in calculating "Commit statements per second."
- The type is int.
- The unit is commits.

#### Efficiency Pages per Prefetch for Interval

- The number of data pages read per prefetch request for the database during the monitoring interval. Use the returned value to determine the amount of asynchronous input and output in each interaction with the prefetcher. An excessively low returned value when compared to normal operating levels indicates that you need more input and output servers. The more input and output servers that you have, the better your query performance.
- The type is int.
- The unit is pages/prefetch.

#### Efficiency Deadlock Rollbacks Percent

- The percentage of the total number of rollbacks that deadlock caused. The value format is an integer.

- The type is double.
- The unit is percent.

#### Efficiency Select SQL Percent for Interval

- The percentage of total SQL statements that were SQL SELECT statements during the monitoring interval. Use the returned value to determine the level of application activity and throughput for the database.
- The type is double.
- The unit is percent.

#### Efficiency Database Efficiency Total Sync I/O Time

- The total time in milliseconds applied to processing requests for synchronous reads or writes for the database. The returned value is the sum of the returned values from the average pool write time (ms) and average pool read time (ms). This time is the amount of time that database agents spend doing synchronous reads and writes.
- The type is int.
- The unit is milliseconds.

#### Efficiency Failed SQL Stmts Percent for Interval

- The percentage of total Structured Query Language statements that failed during the monitoring interval. The value format is an integer.
- The type is double.
- The unit is percent.

#### Efficiency Database Efficiency Pool Sync Index Writes

- The number of physical index write requests that were performed synchronously (that is, physical index page writes that were performed by database manager agents). This value is derived by subtracting the value of the Pool Async Index Writes attribute from the value of the Pool Index Writes attribute. By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing.
- The type is int.
- The unit is writes.

#### Efficiency Database Efficiency Avg Pool I/O Time

- The average time (in milliseconds) for performing buffer pool input and output operations (reading or writing) to the database. A high average time can indicate the existence of an input and output conflict. In this case, you might need to move data to a different device. The returned value includes the time applied to asynchronous input and output operations (which are performed by prefetchers and page cleaners). The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Efficiency Page Cleans for Interval

- The number of times a page cleaner was invoked for the database (for any reason) during the monitoring interval. Use the returned value to determine how often pages are written to disk by the page cleaners of this database. If this value increases over time, you can define more page cleaners. The number of page cleaners is determined by the number of I/O cleaners configured. The following value is also valid:
- The type is int.
- The unit is cleans.

#### Efficiency Database Efficiency Total Sync I/O

- The total synchronous input and output.
- The type is int.
- The unit is occurrences.

#### Efficiency Event Monitors

- The number of event monitors defined in the database. Use the returned value to determine how many event monitors are defined for the database. When you define an event monitor, its definition is stored in the database system catalog table. You can create any number of event monitors. However, the maximum number of event monitors that can be active for a database at any given time is 32. The following value is also valid:
- The type is int.
- The unit is monitors.

#### Efficiency Lock Waits for Interval

- The number of times that applications had to wait for locks in the database during the monitoring interval. Use the returned value as an indication of how much time is applied to waiting for locks during a particular monitoring interval.
- The type is int.
- The unit is waits.

#### Efficiency Database Efficiency Binds Precompiles

- The number of binds and precompiles attempted. Use this attribute to gain insight into the current level of activity within the database manager.
- The type is int.
- The unit is binds.

#### Efficiency Invalid System Packages

- The number of system packages that are currently marked not valid in the database. Use the returned value as an indication of the current number of nonvalid packages owned by the system. A package is marked not valid if it depends on an object (for example, a table) and that object is dropped. The number of packages that are not valid can indicate how many automatic rebinds are necessary in the database. The package is automatically rebound the next time it is accessed, unless it was marked not valid because a trigger was dropped or because the dropped object was not recreated. Use of automatic rebinds can significantly lower performance, and must be minimized where possible.
- The type is int.
- The unit is packages.

#### Efficiency Database Efficiency Appl Section Lookups

- The number of lookups of SQL sections by an application from its SQL work area. This counter indicates how many times the SQL work area was accessed by agents for an application. It is a cumulative total of all lookups on all SQL work heaps for agents working on this application. The following value is also valid:
- The type is int.
- The unit is lookups.

#### Efficiency Log I/O for Interval

- The total amount of log input and output. This amount is the sum of the number of log pages read and the number of log pages written within the monitoring interval. Use the returned value

to determine whether you must move the log to a different device. If this input and output is beyond the capabilities of the current device, you can determine if moving the log (by changing the newlogpath configuration parameter) improves performance. The following value is also valid:

- The type is int.
- The unit is pages.

#### Efficiency Avg Pool Async Data Writes

- The average number of buffer pool asynchronous data writes (data and index) when compared to the total number of pool writes for the database. The value is derived through this formula:  $\text{pool async data writes} / (\text{pool data writes} + \text{pool index writes})$ . Use the returned value to gain insight into how well the page cleaners are working and to refine the num\_iocleaners configuration parameter. If the returned value is low compared to normal operating levels, increase the number of input and output cleaners by increasing the value of the num\_iocleaners parameter. If the returned value is high compared to normal operating levels, you can save system resources by decreasing the number of input and output cleaners (by decreasing the value of the num\_iocleaners parameter).
- The type is int.
- The unit is writes.

#### Efficiency Total Pool Phys I/O

- The total time in milliseconds applied to physical I/O for the database. A high returned value (as compared to the total number of physical buffer pool input and output operations) can indicate the presence of an input and output wait, which in turn can indicate that you must move data to a different device.
- The type is int.
- The unit is milliseconds.

#### Efficiency Internal Rows Deleted

- The number of rows deleted from the database as a result of internal activity. Use the returned value to gain insight into internal activity within the database. If this activity is high compared to normal operating levels, you can evaluate your table design to determine if the referential constraints or triggers that you defined on your database are necessary.
- The type is int.
- The unit is rows.

#### Efficiency Database Efficiency UID SQL Percent for Interval

- The percentage of total SQL statements that were SQL UPDATE, INSERT, and DELETE statements during the monitoring interval. Use the returned value to determine the level of database data change activity.
- The type is double.
- The unit is percent.

#### Efficiency Total Pool Phys Write

- The total time in milliseconds for buffer pool physical writes (including asynchronous writes). The value is derived through this formula:  $\text{pool data writes} + \text{pool index writes}$ . The returned value is used to calculate the average pool write time. This average can indicate the presence of an input and output wait, which in turn can indicate that you must move data to a different device.
- The type is int.
- The unit is milliseconds.



#### Efficiency Database Efficiency Prefetch Wait Time

- The time an application spent waiting for an I/O server (prefetcher) to finish loading pages into the buffer pool. This attribute can be used to experiment with changing the number of I/O servers and the I/O server sizes.
- The type is int.
- The unit is milliseconds.

#### Efficiency Avg Appls

- The value of the average number of active applications.
- The type is int.
- The unit is applications.

#### Efficiency Secondary Log Used Percent

- The percentage of maximum log space used by the secondary log. Use the returned value to show the current dependency on secondary logs. Secondary logs are used when you have circular logging (log retention off) and the primary log files are full.
- The type is double.
- The unit is percent.

#### Efficiency Database Efficiency Avg Sync Write Time

- The average elapsed time used to perform a synchronous write. This value is derived by dividing the value of the Pool Sync Write Time attribute by the value of the Pool Sync Write attribute. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device.
- The type is int.
- The unit is milliseconds.

#### Efficiency Internal Rows Inserted

- The number of rows inserted into a database as a result of internal activity caused by triggers. Use the returned value to gain insight into internal activity within the database. If this activity is high compared to normal operating levels, you can evaluate your design to determine if you can alter it to reduce this activity.
- The type is int.
- The unit is rows.

#### Efficiency Database Efficiency Appl Section Inserts

- The number of inserts of SQL sections by an application from its SQL work area. The working copy of any executable section is stored in a unique SQL work area. This value represents the number of times when a copy was not available and therefore was inserted. The following value is also valid:
- The type is int.
- The unit is inserts.

#### Efficiency DB Capture Prun

- The number of rows in the unit-of-work table. Use the returned value to help you determine whether you need to prune the unit -of-work (UOW) table or the change data (CD) table.
- The type is int.
- The unit is rows.

### Efficiency Database Efficiency Lock Escalation for Interval

- The total number of lock escalations for applications connected to this database during the monitoring interval. Exclusive lock escalations are included in this number. Use the returned value to help you evaluate the settings of the LOCKLIST and MAXLOCKS configuration parameters. Lock escalations can result in a decrease in concurrency between applications connected to a database.
- The type is int.
- The unit is escalations.

### **Component: Apply Program**

The Apply Program attributes provide status information about the Apply Program processes that are configured to run on a database manager server. To collect Apply Program attributes, the Apply Program must be configured successfully. The DB2 agent must be located on the control server to collect Apply Program attributes. The control server is often the same as the target database server in an Apply subscription set.

#### **Dimensions**

##### Apply Program Apply ID

- The subscriber user ID that started the Apply Program.
- The type is string.

##### Apply Program Instance Name

- The name of the monitored DB2 instance.
- The type is string.

##### Apply Program Snapshot Timestamp

- The date and time when the database system monitor information was collected. Use this attribute to help relate data chronologically if you are saving the results in a file or database for ongoing analysis.
- The type is timestamp.

##### Apply Program Apply Status

- Indicates the state of the Apply subscription process for every distinct apply ID in the Apply Program subscription sets.
- The type is int.

##### Apply Program DB Name

- The database name on the Apply control server where the subscription set table is stored.
- The type is string.

##### Apply Program Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string.

##### Apply Program Apply Qualifier

- Uniquely identifies which Apply Program processes this subscription set.
- The type is string. This is a key dimension.

#### **Metrics**

##### Apply Program Total Apply Sub Fail

- The number of subscriptions with the same apply ID that the Apply Program failed to replicate. This number includes only active subscriptions that failed with a status that is equal to -1.
- The type is int.
- The unit is subscriptions.

#### Apply Program Total Apply Sub Lag

- The total number of Apply Program subscriptions that have not completed within their scheduled replication interval.
- The type is int.
- The unit is subscriptions.

### **Component: Diagnostic Log**

Information about log record from a given facility.

#### **Dimensions**

##### Diagnostic DB Name

- The real name of the database for which information is collected. This name was given to the database when it was created. The value format is a simple text string with a maximum of 60 bytes. Use this attribute to identify the specific database to which the data applies.
- The type is string. This is a key dimension.

##### Diagnostic Process Name

- The name of the operating system process that created this message.
- The type is string.

##### Diagnostic MSGID

- The unique message identifier of the message. The msgid is the combination of the message type, message number, and level. For example, ADM7513W.
- The type is string.

##### Diagnostic Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string.

##### Diagnostic Timezone Displacement

- The difference between UTC (Coordinated Universal Time, formerly known as GMT) and local time at the application server.
- The type is int.

##### Diagnostic Level

- The severity level of the record.
- The type is string.

##### Diagnostic Instance Name

- The name of the monitored DB2 instance.
- The type is string.

##### Diagnostic TID

- The numerical identifier of the thread that created this message. The following value is valid:

- The type is int.

#### Diagnostic Function Name

- The name of the function that generated the message.
- The type is string.

#### Diagnostic PID

- The identifier of the operating system process that created this message.
- The type is int.

#### Diagnostic Component Name

- The name of the component that created the message.
- The type is string.

#### Diagnostic Impact

- This attribute qualifies the impact of this message from a user's perspective. This clarifies the impact of the message on the business process DB2 is part of. The following values are valid: NONE, UNLIKELY, POTENTIAL, IMMEDIATE, and CRITICAL.
- The type is string.

#### Diagnostic Function String

- A string provides information about the function that generated the message, including product name, component name, function name, and probe number.
- The type is string.

#### Diagnostic Record Type

- The type of the record.
- The type is string.

#### Diagnostic Message Number

- The numeric message number.
- The type is int.

#### Diagnostic Facility

- A facility is a logical grouping which records relate to.
- The type is string.

#### Diagnostic Partition Num

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If a db partition filter is not specified, data is returned for the current database partition. If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range.
- The type is int. This is a key dimension.

#### Diagnostic Message

- The short description text for this record.
- The type is string.

#### Diagnostic Message Type

- The type of the message.
- The type is string.

#### Metrics

##### Diagnostic Timestamp

- The date and time when the message was created.
- The type is timestamp.
- The unit is unspecified.

#### Component: Log

Information about database configuration parameters that are related to archive logs, the number of archive logs, and the size of archived log path.

#### Dimensions

##### Log Record Operation

- The operation identifier.
- The type is string.

##### Log Record Device Type

- The type identifier of the device that is associated with a logged event.
- The type is string.

##### Log Overflow Log Path

- The location for DB2 to find log files that are needed for a rollforward operation, and to store active log files that are retrieved from the archive. It also gives a location for finding and storing log files that are needed for using db2ReadLog API.
- The type is string.

##### Log Record First Log

- The name of the earliest transaction log that is associated with an event.
- The type is string.

##### Log Record Location

- The full path of the location to store files, such as backup images or load input file, that are associated with logged events.
- The type is string.

##### Log Record Last Log

- The name of the latest transaction log that is associated with an event.
- The type is string.

##### Log Record Instance Name

- The name of the monitored DB2 instance.
- The type is string.

##### Log Backup Pending

- This parameter indicates whether you need to do a full backup of the database before accessing it. This parameter is only on if the database configuration is changed so that the database moves from being nonrecoverable to recoverable (that is, initially both the logretain and userexit parameters were set to NO, and then either one or both of these parameters is set to YES, and the update to the database configuration is accepted).
- The type is int.

#### Log Fail Log Path

- This attribute specifies a path to which the monitored DB2 instance will try to archive log files if the log files cannot be archived to either the primary or the secondary (if set) archive destinations because of a media problem affecting those destinations. This specified path must reference a disk. If there are log files in the directory that is specified by the fail log path attribute, any updates to the fail log path attribute will not take effect immediately. Instead, the update will take effect when all applications disconnect.
- The type is string.

#### Log Record DB Name

- The real name of the database for which information is collected. This name was given to the database when it was created. The value format is a simple text string with a maximum of 60 bytes. Use this attribute to identify the specific database to which the data applies.
- The type is string. This is a key dimension.

#### Log Rollforward Pending

- This attribute indicates whether the monitored DB2 instance is the rollforward pending status.
- The type is int.

#### Log Record Backup ID

- The backup identifier or unique table identifier.
- The type is string.

#### Logging Restore Pending

- This attribute indicates whether a RESTORE PENDING status exists in the database.
- The type is int.

#### Log Database Is Consistent

- This attribute indicates whether the database is in a consistent state.
- The type is int.

#### Log Record DB Alias

- The alias of the database for which information is collected. The value format is a simple text string with a maximum of 60 bytes. Use this attribute to identify the specific database to which the data applies.
- The type is string.

#### Log Record DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If you do not specify a db partition filter, data is returned for either the current database partition (single partition environment) or the aggregated database partitions (multiple partition environment). If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data

collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range, the following values are also valid:

- The type is string. This is a key dimension.

#### Log Record Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string.

#### Log Arch Meth1

- The media type of the primary destination for archived logs.
- The type is string.

#### Log Arch Meth2

- The media type of the secondary destination for archived logs.
- The type is string.

#### Log Path

- This value indicates the current path that is used for logging purposes.
- The type is string.

#### Log Mirror Log Path

- The string that is specified for the mirror log path. The string points to a full qualified path name.
- The type is string.

#### Logging New Log Path

- The current value of the newlogpath configuration parameter. You can use the newlogpath configuration parameter to specify a new location for the log files. The specified path does not become the current log path until both of the following conditions are met:
  - The database is in a consistent state.
  - All users are disconnected from the database.

When the first new connection is made to the database, the database manager moves the logs to this location.

- The type is string.

#### Log Record Object Type

- The identifier for the target object of an operation.
- The type is string.

#### Log Record Operation Type

- The action identifier of an operation.
- The type is string.

#### Log Record Entry Status

- The identifier for the status of an entry in the history file.
- The type is string.

## Metrics

### Logging Log Reads

- The number of log pages that are read from disk by the logger.
- The type is int.
- The unit is pages.

#### Log New Log Path Free Size

- The amount of the free space (in MB) of the file system that is pointed by the new log path attribute.
- The type is int.
- The unit is megabytes.

#### Log Fail Log Path Total size

- The capacity of the file system (in MB) that is pointed to by the fail log path attribute.
- The type is int.
- The unit is megabytes.

#### Logging Sec Log Used Top

- The maximum amount of secondary log space (in bytes) that has been used. Use this attribute with the Secondary Logs Allocated and Total Log Used Top attributes to show the current dependency on secondary logs. If this value is high, you might need larger log files, more primary log files, or more frequent COMMIT statements within your application. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal.
- The type is int.
- The unit is bytes.

#### Log Read Time

- The total elapsed time that the logger spends reading log data from the disk. Use this attribute with the log reads, num log read io, and num log data found in buffer attributes.
- The type is int.
- The unit is seconds.

#### Log Retain

- The value of the log retain enable configuration parameter. The attribute is deprecated in DB2 Version 9. 5, but is used in pre-Version 9. 5 data servers and clients.
- The type is int.
- The unit is unspecified.

#### Log Mirror Log Path Free Size

- The amount of the free space (in MB) of the file system that is pointed by the mirror log path attribute.
- The type is int.
- The unit is megabytes.

#### Log Record Sequence Number

- The sequence number. The following value is valid:
- The type is int.
- The unit is unspecified.

#### Logging Num Log Read IO



- The number of I/O requests that are issued by the logger for reading log data from the disk. Use this attribute with the log reads and log read time attributes to determine if the current disk is adequate for logging.
- The type is int.
- The unit is requests.

#### Log Mirror Log Path Total Size

- The capacity (in MB) of the file system that is pointed by the mirror log path attribute.
- The type is int.
- The unit is megabytes.

#### Log New Log Path Total Size

- The capacity (in MB) of the file system that is pointed by the new log path attribute.
- The type is int.
- The unit is megabytes.

#### Log Sec Log Used Percent

- The percentage of maximum log space used by the secondary log. Use the returned value to show the current dependency on secondary logs. Secondary logs are used when you have circular logging (log retention off) and the primary log files are full.
- The type is double.
- The unit is percent.

#### Log Last Active Log

- The file number of the last active log file. Use this attribute with the first active log and current active log attributes to determine the range of active log files. Knowing the range of active log files helps you determine the disk space required for log files. You can also use this attribute to determine which log files have data to help you identify log files needed for split mirror support.
- The type is int.
- The unit is unspecified.

#### Logging Sec Logs Allocated

- The total number of secondary log files that are currently being used for the database. Use this attribute with the Secondary Log Used Top and Total Log Used Top attributes to show the current dependency on secondary logs. If this value is consistently high, you might need larger log files, more primary log files, or more frequent COMMIT statements within your application.
- The type is int.
- The unit is files.

#### Logging Num Log Data Found in Buffer

- The number of times that an agent reads log data from the buffer. Reading log data from the buffer is preferable to reading from the disk because the latter is slower. Use this attribute with the num log read io attribute to determine if the LOGBUFSZ database configuration parameter needs to be increased.
- The type is int.
- The unit is reads.

#### Logging Total Log Used

- The total log space used in bytes in the database. The value format is an integer. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal.
- The type is int.
- The unit is bytes.

#### Log Record Uniquely Identifies

- The number that uniquely identifies an entry in the history file.
- The type is int.
- The unit is unspecified.

#### Log Current Primary Log Used Percent

- The percentage of the primary log space that are currently used.
- The type is double.
- The unit is percent.

#### Logging Num Log Buffer Full

- The number of times that agents have to wait for log data to write to disk while copying log records into the log buffer. This value is increased per agent per incident. For example, if two agents attempt to copy log data while the buffer is full, this value is increased by two. Use this attribute to determine if the LOGBUFSZ database configuration parameter needs to be increased.
- The type is int.
- The unit is occurrences.

#### Log Fail Log Path Free Size

- The amount of free space (in MB) that is available on the file system that is pointed to by the fail log path attribute.
- The type is int.
- The unit is megabytes.

#### Log Primary

- The number of primary log files.
- The type is int.
- The unit is files.

#### Log Arch Meth2 Free Size

- The amount of free space (in MB) of the secondary destination for archived logs.
- The type is int.
- The unit is megabytes.

#### Log Record Start Timestamp

- The date and time that a logged event started.
- The type is timestamp.
- The unit is unspecified.

#### Log Num Arch Retry

- The number of times that the monitored DB2 instance is to try archiving a log file to the primary or the secondary archive directory before trying to archive log files to the failover directory. This parameter is only used if the failarchpath database configuration parameter is set. If the failarchpath database configuration parameter is not set, DB2 will continuously retry archiving to the primary or the secondary log path.
- The type is int.
- The unit is attempts.

#### Logging Total Log Used Pct

- The percentage of the log space that is in used the database. The value is calculated using the following formula:
- The type is double.
- The unit is percent.

#### Log User Exit

- The value of the user exit enable configuration parameter. The parameter is deprecated in DB2 Version 9. 5, but is used in pre-Version 9. 5 data servers and clients.
- The type is int.
- The unit is unspecified.

#### Log Current Archive Log

- The file number of the log file the DB2 instance is currently archiving. If the DB2 instance is not archiving a log file, the value for this element is SQLM\_LOGFILE\_NUM\_UNKNOWN. Use this attribute to determine if there is a problem archiving log files. Such problems include the following two problems:
- The type is int.
- The unit is fileNumber.

#### Log Arch Retry Delay

- The archive retry delay on error configuration parameter of the monitored DB2 instance. This parameter specifies the number of seconds to wait after a failed archive attempt before trying to archive the log file again. Subsequent retries only takes affect if the value of the numarchretry database configuration parameter is at least 1.
- The type is int.
- The unit is seconds.

#### Logging Num Log Part Page IO

- The number of I/O requests that are issued by the logger for writing partial log data to the disk. Use this attribute with the log writes, log write time, and num log write io attributes to determine if the current disk is adequate for logging.
- The type is int.
- The unit is requests.

#### Logging Log Writes

- The number of log pages written to disk by the logger. Use this attribute with an operating system monitor to quantify the amount of I/O on a device that is attributable to database activity.
- The type is int.
- The unit is pages.

#### Log Path Total Size

- The capacity of the file system (in MB) that is pointed by the log path attribute.
- The type is int.
- The unit is megabytes.

#### Log Path Free Size

- The amount of the free space (in MB) on the file system that is pointed to by the log path attribute.
- The type is int.
- The unit is megabytes.

#### Log First Active Log

- The file number of the first active log file. Use this attribute with the last active log and current active log attributes to determine the range of active log files. Knowing the range of active log files helps you determine the disk space required for log files. You can also use this attribute to determine which log files have data to help you identify log files needed for split mirror support. The following value is valid:  
nFirst active log file.
- The type is int.
- The unit is fileNumber.

#### Log Write Time

- The total elapsed time that the logger spends writing log data to the disk. Use this attribute with the log writes and num log write io attributes to determine whether the current disk is adequate for logging.
- The type is int.
- The unit is seconds.

#### Log Second

- The number of secondary log files.
- The type is int.
- The unit is files.

#### Log Buffer Size (4KB)

- This value specifies the amount (in 4KB) of the database heap to use as a buffer for log records before writing these records to disk. It is important that the log buffer can hold the amount of log space used by an average transaction. Otherwise, logging performance decreases and slows the overall system. The valid format is integer.
- The type is int.
- The unit is 4kilobyteBlocks.

#### Log Record End Timestamp

- The date and time that a logged event ended.
- The type is timestamp.
- The unit is unspecified.

#### Log Overflow Log Path Free Size

- The amount of free space (in MB) of the file system that is specified by the overflow log path attribute.
- The type is int.

- The unit is megabytes.

#### Log Current Secondary Log Used Percent

- The percentage of the secondary log space that are currently used.
- The type is double.
- The unit is percent.

#### Logging Num Log Write IO

- The number of I/O requests that are issued by the logger for writing log data to the disk. Use this attribute with the log writes and log write time attributes to determine if the current disk is adequate for logging.
- The type is int.
- The unit is requests.

#### Log Overflow Log Path Total Size

- The capacity (in MB) of the file system that is specified by the overflow log path attribute.
- The type is int.
- The unit is megabytes.

#### Logging Total Log Available

- The amount of active log space in the database that is not being used by uncommitted transactions (in bytes). Use this element in conjunction with the total log used attribute to determine whether you need to adjust the following configuration parameters of the monitored DB2 instance to avoid running out of log space:
- The type is int.
- The unit is bytes.

#### Log Arch Meth1 Free Size

- The amount of free space (in MB) of the primary destination for archived logs.
- The type is int.
- The unit is megabytes.

#### Log Held By Dirty Pages

- The amount of log (in bytes) corresponding to the difference between the oldest dirty page in the database and the top of the active log. When the snapshot is taken, this value is calculated based on conditions at the time of that snapshot. Use this element to evaluate the effectiveness of page cleaning for older pages in the buffer pool. The following value is valid:
- The type is int.
- The unit is bytes.

#### Log Current Active Log

- The file number of the active log file that the monitored DB2 instance is currently writing. Use this attributes with the first active log and last active log attributes to determine the range of active log files. Knowing the range of active log files helps you determine the disk space required for log files. You can also use this attribute to determine which log files have data to help you identify log files needed for split mirror support. The following value is valid:
- The type is int.
- The unit is fileNumber.

#### Logging Primary Log Used Percent

- The percentage of total log space used by the primary log. Use the returned value to help you evaluate the allocated amount of primary log space and refine the log buffer size, log file size, and primary log configuration parameters. The returned value is valid only if circular logging is used.
- The type is double.
- The unit is percent.

#### Log Arch Meth1 Total Size

- The capacity (in MB) of the primary destination for archived logs.
- The type is int.
- The unit is megabytes.

#### Logging Total Log Used Top

- The maximum amount of total log space (in bytes) that has been used. The value format is an integer. Use this attribute to evaluate the amount of primary log space that is allocated. Comparing the value of this attribute with the amount of primary log space that is allocated can help you evaluate the configuration parameter settings. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal.
- The type is int.
- The unit is bytes.

#### Log Arch Meth2 Total Size

- The capacity (in MB) of the secondary destination for archived logs.
- The type is int.
- The unit is megabytes.

#### Log File Size (4KB)

- This value defines the size (in 4KB) of each primary and secondary log file. The size of these log files limits the number of log records that can be written to them before they become full and a new log file is required.
- The type is int.
- The unit is 4kilobyteBlocks.

#### Log to Redo for Recovery

- The amount of log (in bytes) that will have to be redone for crash recovery. When the snapshot is taken, this value is calculated based on conditions at the time of that snapshot. Larger values indicate longer recovery times after a system crash. If the value seems excessive, check the log held by dirty pages attribute to see if page cleaning needs to be tuned. Also check if there are any long running transactions that need to be terminated.
- The type is int.
- The unit is bytes.

### **Component: DB2 TableSpace**

The Tablespace attributes provide information to monitor page size and usage characteristics for a tablespace.

#### **Dimensions**

Tablespace Name

- The tablespace name of DB2. The Value format is a simple text string with a maximum of 96 characters.
- The type is string. This is a key dimension.

#### Tablespace Container Name

- Indicates the location of the container. Value format is a simple text string with a maximum of 768 characters.
- The type is string.

#### Tablespace Page Size

- Indicates the page size.
- The type is int.

#### Tablespace Prefetch Size

- Indicates the prefetch size. A valid value is an integer.
- The type is int.

#### Tablespace Auto Resize TBSP ID

- The identifier for the tablespace.
- The type is int.

#### Tablespace Usable Pages

- Represents the number of usable pages associated with the database.
- The type is int.

#### Tablespace Auto Resize Last Resize Failed

- Indicates whether the last resize attempted succeeded.
- The type is int.

#### Tablespace Num Containers

- Indicates the number of containers used. The following value is also valid:
- The type is int.

#### Tablespace Status

- The status of the tablespace. DB2 sets this value, which corresponds to the DB2 tablespace\_state element in the DB2 tablespace\_nodeinfo Snapshot logical data group. This element contains a hexadecimal value that indicates the current tablespace state. The externally visible state of a tablespace is composed of the hexadecimal sum of certain state values. For example, if the state is quiesced: EXCLUSIVE and Load pending , the value is 0x0004 + 0x0008, which is 0x000c.
- The type is int.

#### Tablespace Auto Resize Host Name

- The hostname of the machine where the DB2 database is hosted.
- The type is string.

#### Tablespace Auto Resize Node Name

- The node name of the database.
- The type is string.

#### Tablespace Total Pages

- The total number of available pages that are associated with the database.
- The type is int.

#### Tablespace Content Type

- The type of content in a tablespace, such as permanent data, system temporary data, and user temporary data.
- The type is string.

#### Tablespace Auto Resize TBSP Name

- The name of the tablespace. The value format is a simple text string with a maximum of 96 characters.
- The type is string. This is a key dimension.

#### Tablespace Auto Resize Instance Name

- The name of the monitored DB2 instance (Unicode).
- The type is string. This is a key dimension.

#### Tablespace Auto Resize DB Name

- The name of the monitored database.
- The type is string. This is a key dimension.

#### Tablespace Auto Resize Last Resize Time

- Represents the time at which the last tablespace resize was successfully executed.
- The type is string.

#### Tablespace Version

- Indicates the version number of DB2.
- The type is string.

#### Tablespace Extent Size

- Indicates the extent size.
- The type is int.

#### Tablespace Auto Resize DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If a db partition filter is not specified, data is returned for the current database partition. If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range,.
- The type is string. This is a key dimension.

#### Tablespace Object ID

- Represents the identifier for the object.
- The type is int.

#### Tablespace ID



- Represents the identifier for the tablespace.
- The type is int.

#### Tablespace Auto Storage state

- The status of automatic storage for a tablespace in a database.
- The type is string.

#### Tablespace Status Name

- The comma-delimited tablespace state name that corresponds to the tablespace status (TBSP STATUS) attribute. The following table shows the text that is used, depending upon the bit setting of tbsp status:
- The type is string.

#### Tablespace Auto Resize Type

- Tablespace Type of DB2 The automatic storage status of a tablespace.
- The type is string.

#### Tablespace Auto Resize Enabled

- Indicates whether automatic resizing is enabled for the tablespace.
- The type is int.

#### Tablespace Auto Resize Using Auto Storage

- The status of automatic storage for a tablespace in a database.
- The type is int.

#### Tablespace Auto Resize Rebalance Mode

- Represents whether a forward or reverse rebalance is taking place.
- The type is string.

### Metrics

#### Tablespace Direct Reads

- The number of requests to perform a direct read from disk of one or more sectors of data for the tablespace since the first connection. The following value is also valid:
- The type is int.
- The unit is requests.

#### Tablespace Space Used SMS Table

- The number of bytes allocated to the System Managed Space (SMS) tablespace. Use the returned value to determine whether the number of bytes used by the SMS tablespace is excessive in relation to the file system on which the tablespace is established. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal.
- The type is int.
- The unit is bytes.

#### Tablespace Pool I/O per Sec

- The rate (per second) for buffer pool input and output for the tablespace. Buffer pool input and output includes all physical data and index pages that go through the buffer pool when read or written.

- The type is int.
- The unit is poolIO/second.

#### Tablespace Auto Resize Page Size

- Page size.
- The type is int.
- The unit is bytes.

#### Tablespace Pool Async Read Time

- The total time (in milliseconds) that database manager prefetchers spent reading data into the buffer pool for the tablespace. Compare the returned value to the synchronous read time to understand where input and output time is being spent.
- The type is int.
- The unit is milliseconds.

#### Tablespace Pool Read Time

- The time (in milliseconds) spent reading data from the buffer pool to disk for the tablespace since the first connection.
- The type is int.
- The unit is milliseconds.

#### Tablespace Pool Write Time

- The time (in milliseconds) spent writing data from the buffer pool to disk for the tablespace since the first connection.
- The type is int.
- The unit is milliseconds.

#### Tablespace Direct Writes

- The number of direct writes to disk for the tablespace since the first connection. The returned value is used in calculating the returned value for the average number of sectors written per direct write. Direct writes are performed in units, the smallest being a 512-byte sector. They are used while the system is doing any of the following operations: writing LONG VARCHAR columns, writing LOB columns, performing a restore, or performing a load. The following value is also valid:
- The type is int.
- The unit is writes.

#### Tablespace Used Pages

- The total number of used pages.
- The type is int.
- The unit is pages.

#### Tablespace Avg Sector Written

- The average number of sectors that are written for this tablespace for each direct read. Direct writes do not use the buffer pool, and so result in poor performance because the data is physically written from disk each time. If you are using system monitors to track input and output for the device, this returned value helps you distinguish database input and output from non-database input and output. The following value is also valid:
- The type is int.

- The unit is sectors.

#### Tablespace Avg Pool I/O Time

- The average time (in milliseconds) for performing buffer pool input and output operations (reading or writing) for the tablespace. A high average time can indicate the existence of an input and output conflict. In this case, you might need to move data to a different device. The returned value includes the time applied to asynchronous input and output operations (which are performed by prefetchers and page cleaners). The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Tablespace Pool Async Data Reads

- The number of data pages read asynchronously into the buffer pool for the tablespace. Compare the returned value with number of synchronous reads to gain insight into how well the prefetchers are working.
- The type is int.
- The unit is pages.

#### Tablespace Total Sync I/O Time

- The total time (in milliseconds) for processing requests for synchronous reads or writes within the tablespace.
- The type is int.
- The unit is milliseconds.

#### Tablespace Pool Async Data Read Reqs

- The number of asynchronous data read requests.
- The type is int.
- The unit is requests.

#### Tablespace Pool Hit Percent

- The percent buffer pool hit ratio (data plus index).
- The type is double.
- The unit is percent.

#### Tablespace Pool Async Index Read Reqs

- The number of asynchronous index read requests.
- The type is int.
- The unit is requests.

#### Tablespace Avg Sync Data Read Time

- The average time (in milliseconds) for synchronous data reads for the tablespace. Use the returned value to analyze the input and output work being performed for the tablespace. Synchronous read operations are performed by database manager agents. Asynchronous reads are performed by prefetchers, which read data pages from disk into the buffer pool in anticipation of their use. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Tablespace Pool Sync Index Writes

- The number of buffer pool synchronous index writes.
- The type is int.
- The unit is writes.

#### Tablespace Space Used DMS Table Percent

- The percentage of space used in the Database Managed Space (DMS) tablespace. Use the returned value to determine if the tablespace needs more space.
- The type is double.
- The unit is percent.

#### Tablespace Auto Resize Utilization Pct

- The percentage of used pages for a tablespace.
- The type is double.
- The unit is percent.

#### Tablespace Prefetch Percent for Interval

- The percentage of asynchronous read requests that were satisfied for a tablespace during the last monitoring interval.
- The type is double.
- The unit is percent.

#### Tablespace Direct Read Time

- The time (in milliseconds) for performing the direct reads for the tablespace since the first connection. The returned value is used in calculations for the average direct read time (ms). A high average time can indicate an input and output conflict. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Tablespace Auto Resize Used Pages

- The total number of used pages.
- The type is int.
- The unit is pages.

#### Tablespace Pool Index P Reads

- The number of physical read requests to get index pages into the buffer pool for the tablespace.
- The type is int.
- The unit is requests.

#### Tablespace Avg Pool Write Time

- The average time (in milliseconds) for processing write requests that caused data or index pages to be physically written from buffer pool to disk for the tablespace. A high average time generally indicates the existence of an input and output conflict. In this case, you might need to move data to a different device. The returned value includes the time applied to asynchronous write operations that are performed by page cleaners. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Tablespace Total I/O Percent

- The percentage total of I/O.
- The type is double.
- The unit is percent.

#### Tablespace Pool Data P Reads

- The number of read requests requiring input and output to get data pages into the buffer pool for the tablespace since the first connection.
- The type is int.
- The unit is requests.

#### Tablespace Sync Write Time

- The time (in milliseconds) spent synchronously writing data to disk from the buffer pool for the tablespace. Compare the returned value to the value returned by the buffer pool async write time to understand where input and output time for this tablespace is used.
- The type is int.
- The unit is milliseconds.

#### Tablespace Auto Resize Max Size

- Represents the maximum tablespace size in bytes associated with the database.
- The type is int.
- The unit is bytes.

#### Tablespace Avg Sectors Read

- The average number of sectors that are read for this tablespace for each direct read. Direct reads do not use the buffer pool, and so result in poor performance because the data is physically read from disk each time. If you are using system monitors to track input and output for the device, this returned value helps you distinguish database input and output from non-database input and output. The following value is also valid:
- The type is int.
- The unit is sectors.

#### Tablespace Pool Async Write Time

- The total time (in milliseconds) that database manager page cleaners spent writing data or index pages from the buffer pool to disk for the tablespace. Compare the returned value to the synchronous write time to understand where input and output time is being spent.
- The type is int.
- The unit is milliseconds.

#### Tablespace Auto Resize Used/Max Pct

- Represents the percentage of the maximum size that has been used. :
- The type is double.
- The unit is percent.

#### Tablespace Prefetch Reqs for Interval

- The number of prefetch requests for the tablespace during the monitoring interval.
- The type is int.
- The unit is requests.

#### Tablespace Pool Async Index Reads

- The number of index pages read asynchronously into the buffer pool by a prefetcher within the tablespace. By comparing the ratio of asynchronous to synchronous reads, you can determine how well the prefetchers are working. Asynchronous reads are performed by database manager prefetchers.
- The type is int.
- The unit is pages.

#### Tablespace Pool Index from Estore

- The number of buffer pool index pages copied from extended storage for the tablespace.
- The type is int.
- The unit is pages.

#### Tablespace Total Pool I/O Time

- The total pool I/O time.
- The type is int.
- The unit is milliseconds.

#### Tablespace Direct Write Reqs

- The number of requests to perform a direct write to disk of one or more sectors of data for the database. The returned value is used in calculating the returned value for the average number of sectors written per direct write. Direct writes are performed in units, the smallest being a 512-byte sector. They are used while the system is doing any of the following operations: writing LONG VARCHAR columns, writing LOB columns, performing a restore, or performing a load. The following value is also valid:
- The type is int.
- The unit is requests.

#### Tablespace Avg Sync I/O Time

- The average time (in milliseconds) for synchronous input and output operations for the tablespace. Use the returned value to analyze the input and output work being performed for the tablespace. Synchronous input and output operations are performed by database manager agents. Asynchronous input and output operations are performed by prefetchers (reads) and page cleaners (writes). In general, asynchronous input and output helps your applications run faster. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Tablespace Pending Free Pages

- Represents the number of pending free pages associated with the database.
- The type is int.
- The unit is pages.

#### Tablespace Auto Resize Free Pages

- The total number of free pages that are associated with the database.
- The type is int.
- The unit is pages.

#### Tablespace Pool Index L Reads

- The number of logical read requests for index pages that went through the buffer pool for the tablespace since the connection. The returned value includes requests for index pages that are already in the buffer pool or read from disk into the buffer pool to fulfill the request.
- The type is int.
- The unit is requests.

#### Tablespace Total Pool P Write Time

- The total pool physical write time.
- The type is int.
- The unit is milliseconds.

#### Tablespace Total Pool P Read Time

- The time (in milliseconds) spent reading data and index pages from the buffer pool to the disk for the tablespace since the first connection. The value format is an integer. This attribute is the same as the pool read time attribute.
- The type is int.
- The unit is milliseconds.

#### Tablespace Pool Sync Index Reads

- The number of buffer pool synchronous index reads.
- The type is int.
- The unit is reads.

#### Tablespace Auto Resize Total Pages

- The total number of available pages that are associated with the database.
- The type is int.
- The unit is pages.

#### Tablespace Pool Async Index Writes

- The number of times a buffer pool index page was written asynchronously to disk for the tablespace. Subtract the returned value from the buffer pool index writes to calculate the number of synchronous index writes. By comparing the number of asynchronous index writes to synchronous index writes, you can gain insight into how well the buffer pool page cleaners are performing. This ratio can be helpful when you are refining the `num_iocleaners` configuration parameter.
- The type is int.
- The unit is writes.

#### Tablespace Avg Direct Read Time

- The time (in milliseconds) for performing the direct reads for the tablespace. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Tablespace Free Pages

- The number of free pages that are associated with the database.
- The type is int.
- The unit is pages.

#### Tablespace Pool Async Data Writes

- The number of times a buffer pool data page was physically written asynchronously to disk for the tablespace. Compare the returned value with the number of synchronous writes to gain insight into how well the page cleaners are working.
- The type is int.
- The unit is writes.

#### Tablespace Auto Resize Current Size

- Represents the current tablespace size in bytes associated with the database.
- The type is int.
- The unit is bytes.

#### Tablespace Auto Resize Time Stamp

- The local time at the agent when the data was collected.
- The type is timestamp.
- The unit is unspecified.

#### Tablespace Pool Data to Estore

- The number of buffer pool data pages copied to extended storage for the tablespace.
- The type is int.
- The unit is pages.

#### Tablespace Pool Hit Ratio for Interval

- The overall buffer pool hit ratio (as a percentage) for the database during the monitoring interval. This hit ratio includes both index and data page activity. The overall buffer pool hit ratio indicates the percentage of page requests for which the database manager did not need to load a page from disk to service. (That is, the page was already in the buffer pool.) The greater the buffer pool hit ratio, the lower the frequency of disk input and output. If the hit ratio is low compared to normal operating levels, increasing the number of buffer pool pages can improve performance. A ratio of zero indicates that pages needed to be read for every request. For a large database, increasing the buffer pool size can have a minimal effect on the buffer pool hit ratio. Such a database can have so large a number of data pages that the statistical chance of a hit is not increased by an increase of the buffer pools. However, even though the data might be too large to fit in the buffer pool, the entire index can fit. In this case, you can refine buffer pool sizes until the overall buffer pool hit ratio stops increasing, and then refine the buffer pool until the buffer pool index hit ratio no longer increases.
- The type is double.
- The unit is percent.

#### Tablespace Total Sync I/O

- The total number of synchronous reads and writes for both data and index pages for the tablespace.
- The type is int.
- The unit is readsWrites.

#### Tablespace Sync Read Time

- The time (in milliseconds) applied to synchronous reads for the tablespace. Compare the returned value to the buffer pool async read time to understand where input and output time for this tablespace is used.
- The type is int.



- The unit is milliseconds.

#### Tablespace Auto Resize Initial Size

- Represents the initial tablespace size in bytes associated with the database.
- The type is int.
- The unit is bytes.

#### Tablespace Table Space Pool Sync Data Reads

- The number of buffer pool synchronous reads.
- The type is int.
- The unit is reads.

#### Tablespace Pool Data Reads

- The number of read requests to get data pages into the buffer pool for the tablespace.
- The type is int.
- The unit is requests.

#### Tablespace Pool Index Hit Percent for Interval

- The percent buffer pool index hit ratio for the monitoring interval.
- The type is double.
- The unit is percent.

#### Tablespace Auto Resize Usable Pages

- The total number of usable pages that are associated with the database.
- The type is int.
- The unit is pages.

#### Tablespace Total Direct I/O Time

- The total time (in milliseconds) for direct reads to and writes from the tablespace.
- The type is int.
- The unit is milliseconds.

#### Tablespace Avg Sync Data Write Time

- The average time (in milliseconds) for synchronous data writes for the tablespace. Use the returned value to analyze the input and output work being performed for the tablespace. Synchronous write operations are performed by database manager agents. Asynchronous writes are performed by page cleaners, which write out changed pages to disk and free up space in the buffer pool. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Tablespace Direct Read Reqs

- The number of requests to perform a direct read from disk of one or more sectors of data for the database. The returned value is used in calculating the returned value for the average number of sectors read per direct read for the tablespace. Direct reads are performed in units, the smallest being a 512-byte sector. They are used while the system is doing any of the following operations: reading LONG VARCHAR columns, reading LOB columns, or performing a backup. The following value is also valid:
- The type is int.

- The unit is requests.

#### Tablespace Files Closed

- The total number of closed files for the tablespace since the first database connection. The following value is also valid:
- The type is int.
- The unit is files.

#### Tablespace Avg Direct Write Time

- The time (in milliseconds) for performing the direct writes for the tablespace. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Tablespace Auto Resize Increase Size

- Represents the size in bytes of the increase of the tablespace associated with the database.
- The type is int.
- The unit is bytes.

#### Tablespace Pool Data Writes

- The number of times that a buffer pool data page was physically written to disk for the tablespace.
- The type is int.
- The unit is writes.

#### Tablespace Pool Index Writes

- The number of times that a buffer pool index page was physically written to disk for the tablespace since the first connection. If the returned value is high compared to the buffer pool index physical reads, you can improve performance by increasing the available buffer pool space.
- The type is int.
- The unit is writes.

#### Tablespace Table Space Pool Sync Data Writes

- The number of buffer pool synchronous writes.
- The type is int.
- The unit is writes.

#### Tablespace Pool Index to Estore

- The number of buffer pool index pages copied to extended storage within the tablespace. Pages are copied from the buffer pool to extended storage when they are selected as victim pages. This copying is required to make space for new pages in the buffer pool.
- The type is int.
- The unit is pages.

#### Tablespace Auto Resize Used/Total Pct

- Represents the percentage of the allocated size that has been used.
- The type is double.

- The unit is percent.

#### Tablespace Estore Read/Write Ratio

- The ratio (as a percentage) of pages (data plus index) copied from extended storage to pages copied to extended storage within the tablespace. When a page is transferred from extended storage to the buffer pool, you save a system input and output call. However, you still incur the cost of attaching to the extended memory segment, copying the page, and detaching from the segment. Use the returned value to determine if you would benefit from using extended storage. The higher the ratio, the more likely you are to benefit. In general, extended storage is particularly useful if input and output activity is very high on your system.
- The type is double.
- The unit is percent.

#### Tablespace Auto Resize Prefetch Size

- The maximum number of pages that the prefetcher can get from the disk.
- The type is int.
- The unit is pages.

#### Tablespace Avg Pool Read Time

- The average time (in milliseconds) for processing read requests that caused data or index pages to be physically read from disk to buffer pool for the tablespace. A high average time generally indicates the existence of an input and output conflict. In this case, you might need to move data to a different device. The returned value includes the time applied to asynchronous read operations that are performed by prefetchers. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Tablespace Pool Data from Estore

- The number of buffer pool data pages copied from extended storage within the tablespace to the buffer pool. Required pages are copied from extended storage to the buffer pool if they are not in the buffer pool but are in extended storage. This copying can incur the cost of connecting to the shared memory segment but saves the cost of a disk read.
- The type is int.
- The unit is pages.

#### Tablespace Direct Write Time

- The time (in milliseconds) for performing the direct writes for the tablespace since the first connection. The returned value is used in calculations for the average direct write time (ms). A high average time can indicate an input and output conflict. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Tablespace Auto Resize Used/Disk Pct

- Represents the percentage of used tablespace size for the disk size.
- The type is double.
- The unit is percent.

#### Tablespace Pool Data L Reads

- The number of logical read requests for data pages that went through the buffer pool for the tablespace since the connection occurred. The returned value includes requests for data that is already in the buffer pool or read from disk into the buffer pool to fulfill the request.
- The type is int.
- The unit is requests.

#### Tablespace Auto Resize Increase Size Pct

- Represents the size as a percentage to which an auto-resize tablespace increases before reaching its maximum size and more space is required.
- The type is int.
- The unit is bytes.

### **Component: DCS Database**

The DCS Database attributes provide Direct Connection Service (DCS) database information for the monitored database gateway. You can use this information to monitor DCS database specific attributes, such as DCS connection response times and communication errors.

#### **Dimensions**

##### DCS Node Name

- The format is instanceid:hostname:UD for all operating systems. The format for version 6, release 1 of the DB2 Connect agent on Windows systems is instanceid:hostname:UD; on UNIX and Linux systems, the format is instanceid:hostname.
- The type is string.

##### DCS DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If a db partition filter is not specified, data is returned for the current database partition. If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range, the following values are also valid:
- The type is string. This is a key dimension.

##### DCS DB Name

- The real name of the host database for which information is collected or to which the application is connected. This name was given to the database when it was created. The value format is a simple text string with a maximum of 60 bytes.
- The type is string. This is a key dimension.

##### DCS Instance Name

- The name of the monitored DB2 instance (Unicode).
- The type is string.

#### **Metrics**

##### DCS GW Cur Cons

- The current number of connections to host databases that the DB2 Connect gateway is handling. The following value is valid:
- The type is int.
- The unit is connections.

#### DCS Time per Stmt

- The statement execution time (in seconds) divided by the number of attempted statements. The value format is an integer. The following value is valid:
- The type is int.
- The unit is seconds/attempts.

#### DCS Recent Con Rsp Time

- The elapsed time (in seconds) between the start of connection processing and actual establishment of a connection for the most recent DCS application that connected to this database. The following value is valid:
- The type is double.
- The unit is seconds.

#### DCS Network Time per Stmt

- The total host response time minus the total statement execution time divided by the total number of attempted statements. The value format is an integer. The following value is valid:
- The type is int.
- The unit is seconds/attempts.

#### DCS GW Cons Wait Host

- The current number of connections to host databases that the DB2 Connect gateway is handling, and that are waiting for a reply from the host. The following value is valid:
- The type is int.
- The unit is connections.

#### DCS GW Comm Errors for Interval

- The number of times during the monitoring interval that a communication error (SQL30081) occurred while a DCS application tried to connect to a host database, or while it was processing an SQL statement. The following value is valid:
- The type is int.
- The unit is errors.

#### DCS Host Throughput for Interval

- The host throughput in bytes per second for the monitoring interval. This number represents bytes sent plus the number of bytes received divided by the cumulative host response time. The value format is an integer. The following value is valid:
- The type is int.
- The unit is bytes/second.

#### DCS Host Time per Stmt for Interval

- The host response time (in seconds) over the last interval, including any network time over the last interval, divided by the number of statements attempted over the last interval. The value format is an integer. The following value is valid:
- The type is int.
- The unit is seconds.

### **Component: Application**

Information about application activities.

## Dimensions

### Application01 DB Name

- The real name of the database for which information is collected or to which the application is connected. This name was given to the database when it was created. The value format is a simple text string with a maximum of 60 characters. Use this attribute to identify the specific database to which the data applies.
- The type is string.

### Application01 Prev UOW Stop Timestamp

- The date and time that the unit of work completed.
- The type is timestamp.

### Application01 Stmt Start Timestamp

- The date and time that the most recent SQL statement operation started.
- The type is timestamp.

### Application01 Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string.

### Application01 Appl Name

- The name of the application running at the client as it is known to the database manager or DB2 Connect. The value format is a text string, with a maximum of 60 characters. For example: \*Local. db2inst1. 990212202018 .
- The type is string. This is a key dimension.

### Application01 Instance Name

- The name of the monitored DB2 instance.
- The type is string.

### Application01 UOW Comp Status

- The completion status of the previous UOW (unit of work). The value format is a text string with a maximum of 32 characters.
- The type is string.

### Application01 Stmt Stop Timestamp

- The date and time that the most recent SQL statement operation stopped. If the statement is still running, this field is 0 (zero). Use this attribute with the Statement Start attribute to calculate the elapsed execution time for the statement operation.
- The type is timestamp.

### Application01 Appl ID

- The application ID.
- The type is string. This is a key dimension.

### Application01 DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If a db partition filter is not specified, data is returned for the current database partition. If a db partition filter is set to

Aggregated, only aggregated partition data is returned. Historical data collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range, the following values are also valid:

- The type is string. This is a key dimension.

#### Application01 UOW Stop Timestamp

- The date and time that the most recent unit of work completed, which occurs when database changes are committed or rolled back.
- The type is timestamp.

#### Application01 UOW Start Timestamp

- The date and time that the unit of work first required database resources. This resource requirement occurs at the first SQL statement execution for the unit of work.
- The type is timestamp.

#### Application01 Section Number

- The internal section number in the package for the SQL statement currently processing or most recently processed. The value format is an integer.
- The type is int.

### Metrics

#### Application01 Agent Process Agent User CPU Time

- The total CPU time (in seconds) that the database manager agent process spent in system calls. This counter includes time spent on both SQL and non-SQL statements, and any unfenced user-defined functions (UDFs) or stored procedures issued by the application.
- The type is double.
- The unit is seconds.

#### Application01 Agent Sys CPU Time

- The total system CPU time (in seconds) that the database manager agent process spent executing database manager code. This element includes CPU time for both SQL and non-SQL statements, and CPU time for any unfenced user-defined functions (UDFs).
- The type is double.
- The unit is seconds.

#### Application01 Open Curs Blk

- The number of local and remote blocking cursors that are currently open for this application.
- The type is int.
- The unit is cursors.

#### Application01 Avg Sect Read per Direct Read

- The average number of sectors that are read by a direct read for the database. The value is derived through this formula:  $\text{direct reads} / \text{direct read reqs}$ . Direct reads do not use the buffer pool, and so result in poor performance because the data is physically read from disk each time. If you are using system monitors to track input and output for the device, this value helps you distinguish database input and output from non-database input and output.
- The type is int.
- The unit is sectors.

#### Application01 Lock Escalation for Interval

- The total number of lock escalations for the application during the monitoring interval. Exclusive lock escalations are included in this number. Use the returned value to help you evaluate the settings of the LOCKLIST and MAXLOCKS configuration parameters. Lock escalations can result in a decrease in concurrency among the applications connected to a database.
- The type is int.
- The unit is escalations.

#### Application01 Pool Hit Ratio Pct for Interval

- The overall buffer pool hit ratio (as a percentage) for the database during the monitoring interval. This hit ratio includes both index and data page activity. The overall buffer pool hit ratio indicates the percentage of page requests for which the database manager did not need to load a page from disk to service. (That is, the page was already in the buffer pool.) The greater the buffer pool hit ratio, the lower the frequency of disk input and output. If the hit ratio is low compared to normal operating levels, increasing the number of buffer pool pages can improve performance. A ratio of zero indicates that pages needed to be read for every request. For a large database, increasing the buffer pool size can have a minimal effect on the buffer pool hit ratio. Such a database can have so large a number of data pages that the statistical chance of a hit is not increased by an increase of the buffer pools. However, even though the data might be too large to fit in the buffer pool, the entire index can fit. In this case, you can refine buffer pool sizes until the overall buffer pool hit ratio stops increasing, and then refine the buffer pool until the buffer pool index hit ratio no longer increases.
- The type is double.
- The unit is percent.

#### Application01 Total Pool IO Time

- The total time (in seconds) that an application spent performing buffer pool input and output operations (reading or writing pages). The returned value is an indication of how much time the application performs input and output operations using the buffer pool.
- The type is int.
- The unit is seconds.

#### Application01 UID SQL Percent for Interval

- The percentage of total SQL statements that are SQL UPDATE, INSERT, and DELETE statements issued by the application during the monitoring interval. Use the returned value to determine if the application performs frequent updates. If the returned value is low compared to normal operating levels, the application is query-based; otherwise, it is update-based. Knowing what type of applications you have (query-based or update-based) can aid you in refining the database configuration parameters.
- The type is double.
- The unit is percent.

#### Application01 Prefetch Wait Time

- The time an application spent waiting for an I/O server (prefetcher) to finish loading pages into the buffer pool. The value format is an integer. This attribute can be used to experiment with changing the number of I/O servers and the I/O server sizes.
- The type is int.
- The unit is milliseconds.

#### Application01 SQL Reqs Since Commit

- The number of SQL requests that were submitted by the application since the last commit. Use the returned value to monitor the progress of a transaction.



- The type is int.
- The unit is requests.

#### Application01 Deadlocks for Interval

- The total number of deadlocks that occurred for the application during the monitoring interval. Use the returned value to determine if the application is experiencing contention problems. Modify the application to better enable it to run concurrently.
- The type is int.
- The unit is deadlocks.

#### Application01 Lock Wait Time for Interval

- The total elapsed time, in seconds, that the application waited for a lock to be granted during the monitoring interval. The value format is an integer.
- The type is int.
- The unit is seconds.

#### Application01 Avg Sect Written per Direct Write

- The average number of sectors that are written in a direct write by this application.
- The type is int.
- The unit is sectors.

#### Application01 Stmt Sorts

- The total number of times that a set of data was sorted to process the OPEN operation of the current SQL statement. Use the returned value to help identify the need for an index, because indexes can reduce the need for sorting a set of data. Identify the SQL statement for which this returned value is providing sort information. Then, analyze this SQL statement to determine index candidates by looking at columns that are being sorted. For example, a column used in an ORDER BY clause can be an index candidate. The following value is also valid:
- The type is int.
- The unit is sorts.

#### Application01 UOW Log Space Used

- The log space used in the most recent UOW (unit of work). Values that are greater than or equal to 9223372036854775807 are indicated with the text Value Exceeds Maximum in the portal.
- The type is int.
- The unit is bytes.

#### Application01 Pkg Cache Hit Percent

- The application package cache hit ratio (as a percentage) for the last monitoring interval. The package cache hit ratio is the ratio of the difference between the package cache lookups and the package cache inserts to all package cache lookups. This percentage tells you whether the package cache is being used efficiently by this application. If the hit ratio is high (greater than 80%), the package cache is performing well. A smaller percentage can indicate that the package cache must be increased. However, it is not always worthwhile to increase the size of the package cache for an application that runs only once a day. The size of the package cache is set by the pckcachesz configuration parameter.
- The type is double.
- The unit is percent.

#### Application01 Appl Section Inserts

- The number of inserts of SQL sections by an application from its SQL work area. The working copy of any executable section is stored in a unique SQL work area. The returned value is a count of how many times a copy was not available and had to be inserted.
- The type is int.
- The unit is inserts.

#### Application01 Open Curs

- The number of local and remote cursors that are currently open for this application, including the number of local and remote blocking cursors currently open for this application.
- The type is int.
- The unit is cursors.

#### Application01 Appl Work Load

- The ratio of the maximum number of subagents associated with this application to the number of agents that are stolen from the application by DB2 to work on a different application. Use the returned value to evaluate the load that this application places on the system. An agent working for an application is associated with that application. After the agent completes the work for the application, it is placed in the agent pool as an idle agent, but it remains associated with the application. When the application requires an agent again, DB2 searches the agent pool for an agent already associated with the application and assigns work to the associated agent.
- The type is int.
- The unit is ratio.

#### Application01 Lock List in Use Percent

- The percentage of space used in the lock list by a connected application. The value format is a percentage. When an application reaches the maximum number of allowed locks and no additional locks are escalated, the application uses space in the lock list that is allocated for other applications. When an application holds too much of the lock list, other applications can experience lock escalations.
- The type is double.
- The unit is percent.

#### Application01 Total Sorts for Interval

- The total number of sorts that are issued by the application during the monitoring interval. The value format is an integer.
- The type is int.
- The unit is sorts.

#### Application01 Appl Section Lookups

- The number of lookups of SQL sections by an application from its SQL work area. This counter indicates how many times the SQL work area for an application was accessed. The total is a cumulative figure of all lookups on all SQL work heaps for agents working on this application.
- The type is int.
- The unit is lookups.

#### Application01 DL SQL Percent for Interval

- The percentage of total SQL statements that are SQL DDL statements issued by the application during the monitoring interval. Due to the high activity in the system catalog tables, try to keep DDL statement activity to a minimum. If the returned value is high compared to normal operating levels, determine the activity causing it to be high and restrict it from being performed. Examples of DDL statements are CREATE TABLE, CREATE VIEW, ALTER TABLE, and

DROP INDEX. You can also use the returned value to refine the package cache hit ratio for this application. DDL statements can also affect the package cache by invalidating sections that are stored there and causing additional system overhead due to section recompilation.

- The type is double.
- The unit is percent.

#### Application01 Pool Index Hit Ratio Percent for Interval

- The application buffer pool index page hit ratio (as a percentage) during the monitoring interval. The index page hit ratio for the buffer pool indicates the percentage of index page requests for which the database manager did not need to load an index page from disk to service. That is, the index page was already in the buffer pool. The higher the returned value, the lower the frequency of disk input and output, and the faster the performance. If the hit ratio is low compared to normal operating levels, increasing the number of buffer pool pages can improve performance.
- The type is double.
- The unit is percent.

#### Application01 Agents Stolen

- The number of times that agents are stolen from an application. When another application requires a new subagent and has no subagents in its associated agent pool, it steals subagents from the agent pools of other applications. If the number of agents stolen from this application is high compared to normal operating levels, the number of pool agents might be too low. When the agent pool size is too small, one application might fill the pool with associated subagents. When another application requires a new subagent and has no subagents in its associated agent pool, it steals subagents from the agent pools of other applications.
- The type is int.
- The unit is agents.

#### Application01 Associated Agents Top

- The maximum number of associated agents.
- The type is int.
- The unit is agents.

### **Component: High Availability Disaster Recovery**

Information about High Availability Disaster Recovery (HADR) configuration and status.

#### **Dimensions**

##### Disaster Recovery Connect Status

- The current HADR connection status of the monitored database.
- The type is int.

##### Disaster Recovery Standby Log Page

- The page number in the current log file, indicating the current log position on the standby HADR database.
- The type is int.

##### Standby Spool Limit

- The maximum number of pages to spool.
- The type is int.

#### Primary Log File

- The name of the current log file on the primary HADR database.
- The type is string.

#### Disaster Recovery Connect Time

- Depending on the connection status, the value is the HADR connection time, HADR congestion time, or HADR disconnection time.
- The type is timestamp.

#### Disaster Recovery State

- The current HADR state of the database.
- The type is int.

#### Read on Standby Enabled

- Indicates whether the Reads on standby feature is enabled.
- The type is int.

#### Disaster Recovery Standby Log LSN

- The current log position of the standby HADR database. The log sequence number (LSN) is a byte offset in the database log stream.
- The type is int.

#### Primary Instance

- The DB2 instance name on the primary host that is processing the log stream.
- The type is string.

#### Standby Host

- The value of the configuration parameter `hadr_local_host` of the standby member that is processing the log stream.
- The type is string. This is a key dimension.

#### HADR01 Node Name

- Name of origin node of DB2 agent.
- The type is string. This is a key dimension.

#### Disaster Recovery Log Gap

- The average gap (in bytes) between the primary log sequence number (LSN) and the standby LSN.
- The type is int.

#### Disaster Recovery DB Alias

- The alias of the database for which information is collected. The value format is a simple text string with a maximum of 60 bytes. Use this attribute to identify the specific database to which the data applies.
- The type is string.

#### Disaster Recovery Local Host

- The name of the local HADR host. The value is displayed as a host name or an IP address.

- The type is string.

#### Disaster Recovery Database Status

- The status of the monitored database.
- The type is int.

#### Disaster Recovery Peer Window End

- The time until which a HADR primary database stays in the peer or disconnected peer state.
- The type is timestamp.

#### Standby Error Time

- Timestamp of the last error message logged by the standby database.
- The type is timestamp.

#### Disaster Recovery Standby Log File

- The name of the current log file on the standby HADR database.
- The type is string.

#### Peer Wait Limit

- Represents the value of registry variable DB2\_HADR\_PEER\_WAIT\_LIMIT that is used to limit the primary logging wait time in the peer state. The unit is in second.
- The type is int.

#### Disaster Recovery Remote Instance

- The name of the remote HADR instance.
- The type is string.

#### Disaster Recovery Primary Log File

- The name of the current log file on the primary HADR database.
- The type is string.

#### Assisted Member Active

- Returns YES, if the member on primary database that is being assisted is active during assisted remote catchup.
- The type is int.

#### HADR Role

- The current High Availability Disaster Recovery role of the database.
- The type is int.

#### Standby Replay Log File

- The name of the log file corresponding to the standby replay log position on the currently active log stream.
- The type is string.

#### Disaster Recovery DB Location

- The location of the database.
- The type is int.

#### Disaster Recovery Heartbeat

- The number of missed heartbeats on the HADR connection.
- The type is int.

#### Disaster Recovery Primary Log Page

- The page number in the current log file, indicating the current log position on the primary HADR database.
- The type is int.

#### Disaster Recovery Local Service

- The local HADR TCP service. The value is displayed as a service name or a port number.
- The type is string.

#### Primary Host

- The value of the configuration parameter `hadr_local_host` of the member on the primary host that is processing the log stream.
- The type is string.

#### Standby Log File

- The name of the current log file on the standby HADR database.
- The type is string.

#### Disaster Recovery DB Partition

- The mode of the database partition.
- The type is int. This is a key dimension.

#### HADR Syncmode

- The current High Availability Disaster Recovery synchronization mode of the database.
- The type is string.

#### Disaster Recovery Instance Name

- The name of the monitored DB2 instance (Unicode).
- The type is string.

#### Disaster Recovery Remote Host

- The name of the remote HADR host. The value is displayed as a host name or an IP address.
- The type is string.

#### Disaster Recovery Remote Service

- The remote HADR TCP service. The value is displayed as a service name or a port number.
- The type is string.

#### Standby Instance

- The DB2 instance name of the standby member that is processing the log stream.
- The type is string. This is a key dimension.

#### HADR Timeout

- Represents the time period in seconds lapsed, since an HADR database server has confirmed its connection to the partner database is failed, and there is no communication from the partner database.
- The type is int.

#### HADR01 DB Name

- The name of database.
- The type is string. This is a key dimension.

#### Disaster Recovery Primary Log LSN

- The current log position of the primary HADR database. The log sequence number (LSN) is a byte offset in the database log stream.
- The type is int.

#### Peer Window

- Represents a value (in seconds) of `hadr_peer_window`, a configurable parameter of database. This is the configured amount of time for which a HADR primary-standby database pair continues to behave as in a disconnected peer state when the primary database loses connection with the standby database.
- The type is int.

#### HADR Wait Time per Log Flush

- Average log HADR wait time in seconds. Derived as average of Log HADR Wait Time and Log HADR Waits Total.
- The type is int.

#### Query Timestamp

- Date/Time of query execution.
- The type is timestamp.

### Metrics

#### Standby Receive Blocked

- Returns YES, if the standby database temporarily cannot receive logs.
- The type is int.
- The unit is flag.

#### Standby Receive Replay Gap

- The recent average in kilobytes, of the gap between the standby log receive position and the standby log replay position.
- The type is int.
- The unit is kilobytes.

#### Standby Receive Buffer Percent

- Indicates the percentage of standby log receiving buffer that is being used during log shipping. When spooling is enabled, standby can continue to receive logs even when receive buffer is full (that is 100% used).
- The type is double.
- The unit is percent.

#### Standby Log Device Full

- Returns YES, if the standby log device is full.
- The type is int.
- The unit is flag.

#### Standby Key Rotation Error

- Returns YES, if the standby database encountered a master key rotation error.
- The type is int.
- The unit is flag.

#### Standby Tablespace Error

- Returns YES, if a table space of standby database is in an invalid error state and can no longer replay transactions affecting it.
- The type is int.
- The unit is flag.

#### HADR State

- The current High Availability Disaster Recovery state of the database.
- The type is int.
- The unit is state.

#### HADR Log Gap

- Shows the recent average of the gap between the value PRIMARY LOG POS and value STANDBY LOG POS. The gap is measured in number of kilobytes.
- The type is int.
- The unit is kilobytes.

#### Standby Replay Only Window Active

- Indicates whether the DDL or maintenance-operation replay is in progress on the standby.
- The type is int.
- The unit is state.

#### HADR Disconnect Time Left

- Time left to close HADR connection in seconds. Derived from Heartbeat Timeout and Time Since Last Recv.
- The type is int.
- The unit is seconds.

#### Overall HADR Status

- The comprehensive HADR connection status for all partner databases. The status returns 'Critical' when HADR state for primary database or principle standby is DISCONNECTED. It is derived as 'Warning' if the HADR state for auxiliary/secondary standby is DISCONNECTED. Otherwise the peer DB status is 'Normal'.
- The type is int.
- The unit is status.

#### Row Number

- Row Number.



- The type is int.
- The unit is unspecified.

#### HADR Log Delay

- Calculated HADR log delay in seconds. Derived from Primary Log Time and Standby Log Time.
- The type is int.
- The unit is seconds.

#### Standby Replay Not on Preferred

- Returns YES, if the current replay member on the standby is not the preferred replay member.
- The type is int.
- The unit is flag.

#### Standby Spool Percent

- The percentage of spool space used, relative to the configured spool limit.
- The type is double.
- The unit is percent.

#### Disaster Recovery Application Current Connections

- The number of applications that are connected to the database.
- The type is int.
- The unit is connections.

#### Heartbeat Miss Rate

- The rate of missed heartbeats. It is derived from Heartbeat Expected and Heartbeat Missed.
- The type is double.
- The unit is percent.

### **Component: Slow SQL Statements**

Information about slow SQL Statements.

#### **Dimensions**

##### Slow SQL DB Partition

- The DB2 database partition node number. DB2 partition numbers range from 0 to 999. The 'Aggregated' and 'Current' values can be used within a query or situation filter. If no db partition filter is specified, then a row of data will be returned for each database partition. If a db partition filter is used with the 'Aggregated' value, then only aggregated partition data will be returned. Historical data collection will include both aggregated and individual partition attribute data.
- The type is int. This is a key dimension.

##### Slow SQL Instance Name

- The name of the monitored DB2 instance.
- The type is string. This is a key dimension.

##### Slow SQL Node Name

- The origin node of db2 agent.
- The type is string.

#### Slow SQL Statement Type

- The type of the SQL statement, such as static or dynamic.
- The type is string.

#### Slow SQL Statement Text

- The query for the SQL statement.
- The type is string. This is a key dimension.

#### Slow SQL Active State

- The state of the SQL statement.
- The type is string.

#### Slow SQL Duration

- The duration of executing the SQL statement.
- The type is string.

#### Slow SQL Stmt Start Timestamp

- The start time of the SQL statement.
- The type is timestamp.

#### Slow SQL DB Name

- The name of the monitored database.
- The type is string. This is a key dimension.

#### Slow SQL Executable ID

- The unique identifier for the SQL statement.
- The type is string. This is a key dimension.

### **Metrics**

#### Slow SQL Lock wait

- The total number of times that applications or connections waited for locks while executing the SQL Statement.
- The type is int.
- The unit is waits.

### **Component: Custom SQL Execution**

Information about the results of customized SQL statement executions, including five string columns, five number columns, and two date and time columns.

#### **Dimensions**

##### Custom SQL Fourth String Column Name

- The name of the fourth string type column in the result of the customized SQL statement execution.
- The type is string.

##### Custom SQL Second Date Value

- The second date time value in the result of the customized SQL statement execution.
- The type is timestamp.

#### Custom SQL Fifth String Column Name

- The name of the fifth string type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Second String Column Name

- The name of the second string type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Third String Column Name

- The name of the third string type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL First String Column Name

- The name of the first string type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Fifth Number Column Name

- The name of the fifth number type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Fourth Number Column Name

- The name of the fourth number type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Third Number Column Name

- The name of the third number type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Second Number Column Name

- The name of the second number type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Status SQL ID

- The SQL ID that is defined in the definition file.
- The type is string. This is a key dimension.

#### Custom SQL First Number Column Name

- The name of the first number type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Status Node Name

- The managed system name of the agent. For new installations of version 7. 1, the format is instanceid:hostname:UD for all operating systems.
- The type is string.

#### Custom SQL Status Last Execution Error Message

- The error message returned by DB2 for the last SQL execution, which has a maximum length of 256 characters.
- The type is string.

#### Custom SQL Status Status Last Execution Time

- The last date and time when the SQL is executed.
- The type is timestamp.

#### Custom SQL Fifth String Value

- The fifth string value in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Fourth String Value

- The fourth string value in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Status Instance Name

- The name of the monitored DB2 instance.
- The type is string. This is a key dimension.

#### Custom SQL Third String Value

- The third string value in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Second String Value

- The second string value in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL First String Value

- The first string value in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Status DB Alias

- The alias name of the database on which the SQL Statement associated with the SQL ID is executed, which is a key attribute.
- The type is string. This is a key dimension.

#### Custom SQL First Date Column Name

- The name of the first date time type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Status SQL State

- The SQL STATE returned by DB2 for the last SQL execution, which has a length of 10 characters.
- The type is string.

#### Custom SQL DB Alias Filter Name

- The Database alias filter name that can be defined as:
  - The character ( \* ), is required if you want to execute the SQL statement associated with the SQL ID on all the databases of the DB2 server excluding all HADR standby databases.
  - A database alias, this is required if you want to execute the SQL statement associated with the SQL ID on a specific database.

If the database filter alias name contains blank spaces at the beginning and at the end, the blank spaces at the end are trimmed.

- The type is string.

#### Custom SQL Second Date Column Name

- The name of the second date time type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL First Date Value

- The first date time value in the result of the customized SQL statement execution.
- The type is timestamp.

### **Metrics**

#### Custom SQL First Number Value

- The first number value in the result of the customized SQL statement execution.
- The type is int.
- The unit is unspecified.

#### Custom SQL Second Number Value

- The second number value in the result of the customized SQL statement execution.
- The type is int.
- The unit is unspecified.

#### Custom SQL Third Number Value

- The third number value in the result of the customized SQL statement execution.
- The type is int.
- The unit is unspecified.

#### Custom SQL Fourth Number Value

- The fourth number value in the result of the customized SQL statement execution.
- The type is int.
- The unit is unspecified.

#### Custom SQL Fifth Number Value

- The fifth number value in the result of the customized SQL statement execution.
- The type is int.

- The unit is unspecified.

Custom SQL Status Last Execution Error Code

- The native error code returned by DB2 for the last SQL execution.
- The type is int.
- The unit is unspecified.

### **Component: Apply Program Subscription**

The Apply Subscription attributes provide information about Apply Program subscription sets that are configured to run on a database manager server. To collect Apply Program attributes successfully, the Apply Program must be configured. The DB2 agent must be located on the control server to collect Apply Program attributes. The control server is often the same as the target database server in an Apply subscription set.

#### **Dimensions**

Apply Program Subscription Target Owner

- The name of the target owner for this member.
- The type is string. This is a key dimension.

Apply Program Subscription Apply ID

- Subscriber user ID that started the Apply Program.
- The type is string.

Apply Program Subscription Target Table

- The name of the target table or view for this member.
- The type is string. This is a key dimension.

Apply Program Subscription Apply Sub Status

- The Apply Program subscription status.
- The type is int.

Apply Program Subscription Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string.

Apply Program Subscription Instance Name

- The name of the monitored DB2 instance.
- The type is string.

Apply Program Subscription DB Name Target

- Database name.
- The type is string. This is a key dimension.

#### **Metrics**

Apply Program Subscription Apply Sub Lag Time

- The difference (in number of minutes) between how much time has elapsed since the last run of the Apply Program and the expected sleep interval between executions of the Apply Program for the target table. The following value is valid:
- The type is int.

- The unit is minutes.

#### Apply Program Subscription Apply Num Reqs Refresh

- Indicates the number of subscriptions the Apply Program failed to replicate because refresh copying was disabled. While attempting to perform a full refresh, the Apply Program encountered a DISABLE\_REFRESH column in the register table, which was set to On . You can either turn off the DISABLE\_REFRESH column or bypass the Apply Program and perform a manual refresh.
- The type is int.
- The unit is subscriptions.

### Component: Application Activities

Information about application activities.

#### Dimensions

##### Lock Conflict00 Tablespace Name (Unicode)

- The name of the tablespace against which the application currently holds a lock (Unicode). The value format is a text string with a maximum of 60 bytes.
- The type is string.

##### Application Group00 Stmt Start

- The string date and time that the most recent SQL statement operation started. The value format is CYYMMDDHHMMSSmmm. Use this attribute with the Statement Stop attribute to calculate the elapsed execution time for the statement operation.
- The type is string.

##### Lock Conflict00 Client DB Alias

- The alias defined within the database manager where the database connection request originated. The value format is a text string with a maximum of 20 characters. Use to identify the actual database that the application is accessing. The mapping between this name and Database Name can be done by using the database directories at the client node and the database manager server node. Because different database aliases can have different authentication types, this attribute can also help you determine the authentication type.
- The type is string.

##### Application Lock Wait Start Time

- The date and time that the application started waiting to obtain a lock on the object that is currently locked by another application.
- The type is timestamp.

##### Application Group00 Agent Sys CPU Time

- The total system CPU time (in seconds) that the database manager agent process spent executing database manager code. This element includes CPU time for both SQL and non-SQL statements, and CPU time for any unfenced user-defined functions (UDFs).
- The type is string.

##### Application Auth ID

- The authorization ID of the user who invoked the application that is being monitored. On a DB2 Connect gateway node, this ID is the user authorization ID on the host. The value format is a text string with a maximum of 60 bytes. Use this attribute to determine who invoked the application.
- The type is string.

#### Lock Conflict00 Lock Escalation

- An indicator of whether a lock request was made as part of a lock escalation.
  - No
  - Yes

Use this attribute to better understand the cause of deadlocks. If deadlocks occur that involve applications doing lock escalation, you might want to increase the amount of lock memory or change the percentage of locks that any one application can request.

- The type is string.

#### Application Group00 Stmt Stop

- The string date and time that the most recent SQL statement operation stopped. If the statement is still running, this field is 0 (zero). Use this attribute with the Statement Start attribute to calculate the elapsed execution time for the statement operation.
- The type is string.

#### Lock Conflict00 Client DB Alias (Unicode)

- The alias defined within the database manager where the database connection request originated (Unicode). The value format is a text string with a maximum of 60 bytes.
- The type is string.

#### Application Stmt Text

- The text of the dynamic SQL statement. For application snapshots, the statement text helps you identify what the application was executing when the snapshot was taken, or most recently processed if no statement was being processed at the time the snapshot was taken. For dynamic SQL statements, this attribute identifies the SQL text associated with a package. The value format is a text string with a maximum of 2000 bytes.
- The type is string.

#### Application Agent ID Holding Lock

- The application handle of the agent holding a lock for which this application is waiting. The value format is an integer. The lock monitor group must be turned on to obtain this information.
- The type is int.

#### Lock Conflict00 Appl ID (Unicode)

- The identifier generated when the application connects to the database at the database manager or when DDCS receives a request to connect to a DRDA database (Unicode).
- The type is string.

#### Application Table Schema

- The schema of the table the application is waiting to lock. The value format is a text string with a maximum of 60 bytes. Along with the Table Name attribute, this attribute can help to determine the source of contention for resources.
- The type is string.

#### Application Creator

- The authorization ID of the user that precompiled the application (Unicode). The value format is a text string with a maximum of 60 bytes. Use this attribute to help identify the SQL statement that is processing, with the CREATOR column of the package section information in the catalogs.
- The type is string.



#### Application Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string.

#### Application Client Protocol

- The communication protocol that the client application is using to communicate with the server. The value format is a text string with a maximum of 12 characters. Use this attribute for troubleshooting of remote applications.
- The type is string.

#### Application Group00 UOW Stop Time

- The string date and time that the most recent unit of work completed, which occurs when database changes are committed or rolled back. The value format is CYYMMDDHHMMss.
- The type is string.

#### Lock Conflict00 Lock Wait Start Timestamp

- The date and time that the application started waiting to obtain a lock on the object that is currently locked by another application.
- The type is timestamp.

#### Application Appl ID Holding Lock

- The application ID of the application that is holding a lock on the object that this application is waiting to obtain (Unicode). The value format is a text string with a maximum of 96 bytes.
- The type is string.

#### Application Conn Complete Timestamp

- The date and time that a connection request was granted.
- The type is timestamp.

#### Application Client Prdid

- The product and version identifier for the software on the client. The value format is a text string with a maximum of 20 characters. For example: SQL06010 .
- The type is string.

#### Application Appl ID

- The identifier generated when the application connects to the database at the database manager or when DDCS receives a request to connect to a DRDA database. The value format is a text string, with a maximum of 32 characters.
- The type is string. This is a key dimension.

#### Application Client Platform

- The operating system on which the client application is running. Use this attribute to analyze problems for remote applications. The value format is a text string with a maximum of 20 characters.
- The type is string.

#### Application Instance Name

- The name of the monitored DB2 instance.
- The type is string. This is a key dimension.

#### Application Lock Mode

- The type of lock being held. Use the lock mode to determine the source of contention for resources. The value format is a text string with a maximum of 32 characters.
- The type is string.

#### Application Group00 UOW Comp Status

- The completion status of the previous UOW (unit of work). Use this attribute to determine if the unit of work ended due to a deadlock or an abnormal termination.
- The type is string.

#### Lock Conflict00 Snapshot Time

- The string date and time when the database system monitor information was collected. Use this attribute to help relate data chronologically if you are saving the results in a file or database for ongoing analysis. The timestamp value is formatted as a date and time string. The internal timestamp value that is stored in the database is in the format cYYMMDDhhmmss000.
- The type is string.

#### Application Lock Object Type

- The type of object against which the application holds a lock (for object-lock-level information), or the type of object for which the application is waiting to obtain a lock (for application-level and deadlock-level information). The value format is a text string with a maximum of 16 characters.
- The type is string.

#### Application Corr Token

- The DRDA AS correlation token. The value format is a text string with a maximum of 96 bytes.
- The type is string.

#### Application Execution ID

- The ID that the user specified when logging in to the operating system. This ID is distinct from the Authorization ID, which the user specifies when connecting to the database. The value format is a text string with a maximum of 60 bytes. Use this attribute to determine the operating system user ID of the individual running the monitored application.
- The type is string.

#### Lock Conflict00 Auth ID (Unicode)

- The authorization ID of the user who invoked the application that is being monitored (Unicode). On a DB2 Connect gateway node, this is the user authorization ID on the host. The value format is a text string with a maximum of 20 bytes.
- The type is string.

#### Application Group00 Database Manager Agent User CPU Time

- The total CPU time (in microseconds) that the database manager agent process used. This counter includes time spent on both SQL and non-SQL statements, and any fenced user-defined functions (UDFs) or stored procedures issued by the application. System CPU represents the time spent in system calls. User CPU represents time spent executing database manager code. The value format is a text string with a maximum of 10 characters. Use this attribute with the other CPU-time related attributes to help you identify applications or queries that consume large amounts of CPU time.
- The type is string.

#### Application Stmt Type

- The type of SQL statement processed. The value format is a text string with a maximum of 32 characters.
- The type is string.

#### Application Client PID

- The process ID of the client application that made the connection to the database. The value format is an integer. Use this attribute to correlate monitor information such as CPU and I/O time to your client application. If a DRDA AS connection is used, this element is set to 0.
- The type is int.

#### Application DB Name

- The real name of the database for which information is collected or to which the application is connected. This name was given to the database when it was created. The value format is a simple text string with a maximum of 60 bytes. Use this attribute to identify the specific database to which the data applies.
- The type is string.

#### Application Group00 Prev UOW Stop Time

- The string date and time that the unit of work completed. The value format is CYYMMDDHHMMSSmmm. Use this attribute with the UOW Stop Time attribute to calculate the total elapsed time between COMMIT/ROLLBACK points, and with the UOW Start Time attribute to calculate the time spent in the application between units of work.
- The type is string.

#### Lock Conflict00 Table Schema (Unicode)

- The schema of the table against which the application is holding a lock (Unicode). The value format is a text string with a maximum of 60 bytes.
- The type is string.

#### Lock Conflict00 Table Name (Unicode)

- The name of the table against which the application is holding locks (Unicode). The value format is a text string with a maximum of 60 bytes.
- The type is string.

#### Application Appl Status

- The status of the application being monitored. This attribute can help you diagnose potential application problems. The value format is a text string with a maximum of 64 characters.
- The type is string.

#### Lock Conflict00 Status Change Time

- The string date and time the application entered its current status. The value format is CYYMMDDHHMMSSmmm. Use this attribute to determine how long an application has been in its current status. If the application status remains unchanged for a long period of time, the application might have a problem.
- The type is string.

#### Lock Conflict00 Appl Name (Unicode)

- The name of the application running at the client as it is known to the database manager or DB2 Connect (Unicode). The value format is a text string, with a maximum of 60 bytes.

- The type is string. This is a key dimension.

#### Application Package Name

- The name of the package that contains the SQL statement currently executing. The value format is a text string with a maximum of 60 bytes.
- The type is string.

#### Application Table Name

- The name of the table the application is waiting to lock. The value format is a text string with a maximum of 60 bytes. Use this attribute with the Table Schema attribute to determine the source of contention for resources.
- The type is string.

#### Application Group00 Section Number

- The internal section number in the package for the SQL statement currently processing or most recently processed.
- The type is int.

#### Application Group00 UOW Start Time

- The string date and time that the unit of work first required database resources. This resource requirement occurs at the first SQL statement execution for the unit of work. The value format is CYYMMDDHHMMss. Use this attribute with the UOW Stop Time attribute to calculate the total elapsed time of the unit of work and with the Previous Unit of Work Completion Timestamp attribute to calculate the time spent in the application between units of work.
- The type is string.

#### Application Agent ID

- The application handle, which is a system-wide unique ID for the application. The value format is an integer. On multi-node systems, where a database is partitioned, this ID is the same on every node where the application might make a secondary connection.
- The type is int.

#### Application Tablespace Name

- The name of the tablespace that the application is waiting to lock. The value format is a text string with a maximum of 60 bytes. This attribute can help you to determine the source of contention for resources.
- The type is string.

#### Lock Conflict00 Appl ID Holding Lock (Unicode)

- The application ID of the application that is holding a lock on the object that this application is waiting to obtain (Unicode). The value format is a text string with a maximum of 96 bytes.
- The type is string.

#### Application Country Code

- The country code of the client application. The value format is an integer.
- The type is int.

#### Application Appl Name

- The name of the application that is connected to the database. The value format is a text string, with a maximum of 60 bytes. For example: \*Local. db2inst1. 990212202018 .

- The type is string. This is a key dimension.

#### Application DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If a db partition filter is not specified, data is returned for the current database partition. If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range.
- The type is string. This is a key dimension.

#### Application Appl Conn Timestamp

- The date and time that an application started a connection request.
- The type is timestamp.

#### Lock Conflict00 Status Change Timestamp

- The date and time the application entered its current status.
- The type is timestamp.

#### Application Cursor Name

- The name of the cursor corresponding to this SQL statement. The value format is a text string with a maximum of 60 bytes.
- The type is string.

#### Application Stmt Operation

- The statement operation currently being processed or most recently processed (if none is currently running). The value format is a text string with a maximum of 20 characters.
- The type is string.

#### Lock Conflict00 Codepage ID

- The codepage or CCSID at the node where the application started. For snapshot monitor data, this is the code page at the node where the monitored application started. Use this attribute to analyze problems for remote applications. By using this information, you can ensure that data conversion is supported between the application code page and the database code page (or for DRDA host databases, the host CCSID).
- The type is int.

### Metrics

#### Application Pool Data to Estore

- Number of buffer pool data pages copied to extended storage. The value format is an integer.
- The type is int.
- The unit is pages.

#### Application Rows Read

- The number of rows read from the table. The value format is an integer. This attribute helps to identify tables with heavy usage for which you might want to create additional indexes.
- The type is int.
- The unit is reads.

#### Application Direct Writes

- The number of write operations that do not use the buffer pool. The value format is an integer.
- The type is int.
- The unit is writes.

#### Application Group00 Lock Waits

- The total number of times that applications or connections waited for locks. At the database level, the lock waits value is the total number of times that applications waited for locks within this database. At the application-connection level, the lock waits value is the total number of times that this connection requested a lock but waited because another connection was already holding a lock on the data. Use this attribute with the Lock Wait Time attribute to calculate, at the database level, the average wait time for a lock. This calculation can be performed at either the database or the application-connection level. If the average lock wait time is high, look for applications that hold many locks, or have lock escalations, with a focus on tuning your applications to improve concurrency, if appropriate. If escalations are the reason for a high average lock wait time, the values of one or both of the LOCKLIST and MAXLOCKS configuration parameters might be too low.
- The type is int.
- The unit is waits.

#### Application Group00 Locks Held

- The number of locks currently held. If the monitor information is at the database level, this number represents the total number of locks currently held by all applications in the database. If the information is at the application level, this number represents the total number of locks currently held by all agents for the application.
- The type is int.
- The unit is locks.

#### Application Pool Hit Ratio

- The buffer pool hit ratio (as a percentage). The value format is an integer. The sum of the Pool Data Logical Reads and Pool Index Logical Reads attributes is divided by the value of the Pool Total Reads attribute to derive the percentage.
- The type is double.
- The unit is percent.

#### Application Pool Data P Reads

- The number of read requests that required I/O to get data pages into the buffer pool. The value format is an integer.
- The type is int.
- The unit is requests.

#### Application Avg Pool Write Time

- The average elapsed time for a write request. The value format is an integer.
- The type is int.
- The unit is milliseconds.

#### Application UOW Lock Wait Time

- The time the UOW (unit of work) waited on locks (in seconds).
- The type is int.
- The unit is seconds.

#### Application Direct Read Time

- The elapsed time (in milliseconds) required to perform the direct reads. The value format is an integer.
- The type is int.
- The unit is milliseconds.

#### Application Rows Selected

- The number of rows that have been selected and returned to the application. The value format is an integer. Use this attribute to gain insight into the current level of activity within the database manager.
- The type is int.
- The unit is selects.

#### Application Cat Cache Overflows

- The number of times that an insert into the catalog cache failed because the catalog cache was full. The value format is an integer. If the catalog cache overflows value is large, the catalog cache might be too small for the workload. Increasing the size of the catalog cache might improve its performance. If the workload includes transactions that compile a large number of SQL statements referencing many tables, views, and aliases in a single unit of work, compiling fewer SQL statements in a single transaction might improve the performance of the catalog cache. Or if the workload includes the binding of packages containing many SQL statements referencing many tables, views or aliases, you might want to split the packages so that they include fewer SQL statements to improve performance.
- The type is int.
- The unit is overflows.

#### Application Cat Cache Inserts

- The number of times that the system tried to insert table descriptor information into the catalog cache. The value format is an integer. Table descriptor information is inserted into the cache following a failed lookup to the catalog cache while processing a table, view, or alias reference in an SQL statement. The catalog cache inserts value includes attempts to insert table descriptor information that fail due to catalog cache overflow and heap full conditions.
- The type is int.
- The unit is inserts.

#### Application Rows Deleted

- The number of row deletions attempted. The value format is an integer. Use this attribute to gain insight into the current level of activity within the database manager.
- The type is int.
- The unit is deletes.

#### Application Pkg Cache Lookups

- The number of times that an application looked for a section or package in the package cache. The value format is an integer. At a database level, it indicates the overall number of references since the database was started, or monitor data was reset. Note that this counter includes the cases where the section is already loaded in the cache and when the section has to be loaded into the cache.
- The type is int.
- The unit is lookups.

#### Application Int Rollbacks

- The total number of rollbacks initiated internally by the database manager. The value format is an integer.
- The type is int.
- The unit is rollbacks.

#### Application Binds Precompiles

- The number of binds and precompiles attempted. The value format is an integer. Use this attribute to gain insight into the current level of activity within the database manager.
- The type is int.
- The unit is operations.

#### Application Lock Escals

- The number of times that locks have been escalated from several row locks to a table lock. A lock is escalated when the total number of locks held by an application reaches the maximum amount of lock list space available to the application, or the lock list space consumed by all applications is approaching the total lock list space. This data item includes a count of all lock escalations, including exclusive lock escalations. When an application reaches the maximum number of locks allowed and there are no more locks to escalate, the application uses space in the lock list that is allocated for other applications. When the entire lock list is full, an error occurs. The value format is an integer.
- The type is int.
- The unit is occurrences.

#### Application application Static SQL Stmts

- The number of static SQL statements that were attempted. The value format is an integer.
- The type is int.
- The unit is statements.

#### Application Open Local Curs Blk

- The number of local blocking cursors currently open for this application. The value format is an integer.
- The type is int.
- The unit is cursors.

#### Application Commit SQL Stmts

- The total number of SQL COMMIT statements that have been attempted. The value format is an integer. A small rate of change in this counter during the monitor period might indicate that applications are not doing frequent commits. The lack of frequent commits can lead to problems with logging and data concurrency. The following value is also valid:
- The type is int.
- The unit is commits.

#### Application Pool Data from Estore

- Number of buffer pool data pages copied from extended storage. The value format is an integer.
- The type is int.
- The unit is pages.

#### Application Avg Pool Read Time

- The average elapsed time for a read request. The value format is an integer.



- The type is int.
- The unit is milliseconds.

#### Application Pkg Cache Hit Ratio

- The percentage of package sections that were found in the cache. The value format is an integer.
- The type is double.
- The unit is percent.

#### Application Open Local Curs

- The number of local cursors currently open for this application, including those cursors counted by Open Local Cursors with Blocking attribute. The value format is an integer.
- The type is int.
- The unit is cursors.

#### Application Hash Join Small Overflows

- The number of times that hash join data exceeded the available sort heap space by less than 10%. The value format is an integer. If this value and the value of the Hash Join Overflows attribute are high, you must consider increasing the sort heap threshold. If this value is greater than 10% of Hash Join Overflows, you must consider increasing the sort heap size.
- The type is int.
- The unit is occurrences.

#### Application X Lock Escals

- The number of times that locks have been escalated from several row locks to one exclusive table lock, or the number of times an exclusive lock on a row caused the table lock to become an exclusive lock. The value format is an integer. A lock is escalated when the total number of locks held by an application reaches the maximum amount of lock list space available to the application. The amount of lock list space available is determined by the LOCKLIST and MAXLOCKS configuration parameters. Other applications cannot access data held by an exclusive lock.
- The type is int.
- The unit is escalations.

#### Application Hash Join Overflows

- The number of times that hash join data exceeded the available sort heap space. The value format is an integer.
- The type is int.
- The unit is occurrences.

#### Application Avg Sort Time

- The average time that was elapsed to complete a sort operation. At the database or application level, the value for this attribute can indicate the performance issues with sorting. This attribute value is affected by the system load. The following value is also valid:
- The type is int.
- The unit is value.

#### Lock Conflict00 Locks Held

- The number of locks currently held. If the monitor information is at the database level, this is the total number of locks currently held by all applications in the database. If it is at the application

level, this is the total number of locks currently held by all agents for the application. Usage of this attribute depends on the level of information being returned from the database system monitor. The following value is also valid:

- The type is int.
- The unit is locks.

#### Application Select SQL Stmts

- The number of SQL SELECT statements that were issued. The value format is an integer.
- The type is int.
- The unit is selects.

#### Application Pool Index Writes

- The number of times a buffer pool index page was physically written to disk. The value format is an integer. If a buffer pool index page is written to disk for a high percentage of the Pool Index Physical Reads, performance might improve by increasing the number of buffer pool pages available for the database. If all applications are updating the database, increasing the size of the buffer pool might have minimal impact on performance; most pages contain updated data that must be written to disk.
- The type is int.
- The unit is writes.

#### Application Int Rows Inserted

- The number of rows inserted into the database as a result of internal activity caused by triggers. The value format is an integer. This attribute can help to gain insight into the internal activity within the database manager. If this activity is high, you must evaluate the design to determine if you can alter it to reduce this activity.
- The type is int.
- The unit is rows.

#### Application Int Auto Rebinds

- The number of automatic rebinds (or recompiles) that have been attempted. The value format is an integer. Automatic rebinds are the internal binds the system performs when a package has been invalidated. Use this attribute to determine the level of database activity at the application or database level.
- The type is int.
- The unit is rebinds.

#### Application Rows Written

- The number of rows changed (inserted, deleted, or updated) in the table. The value format is an integer. A high value for table-level information indicates heavy usage of the table. If so, you might want to use the Run Statistics (RUNSTATS) utility to maintain efficiency of the packages used for this table.
- The type is int.
- The unit is changes.

#### Application Cat Cache Hit Ratio

- The percentage of catalog sections that are found in the cache. The value format is an integer.
- The type is double.
- The unit is percent.

#### Application Pool Index L Reads

- The number of logical read requests for index pages that have gone through the buffer pool. The value format is an integer.
- The type is int.
- The unit is requests.

#### Application Cat Cache Heap Full

- The number of times that an insert into the catalog cache failed because of a heap full condition in the database heap. The value format is an integer. The catalog cache draws its storage dynamically from the database heap. Even if the cache storage has not reached its limit, inserts into the catalog cache might fail due to a lack of space in the database heap. If the catalog cache heap full count is not zero, you can correct the insert failure condition by increasing the database heap size or by reducing the catalog cache size.
- The type is int.
- The unit is failures.

#### Application Sort Overflows

- The total number of sorts that ran out of sort heap space and might have required disk space for temporary storage. The value format is an integer. at the database or application level, use this element with the Total Sorts attribute. This attribute can help to determine the source of contention for resources.
- The type is int.
- The unit is sorts.

#### Application Cat Cache Lookups

- The number of times that the catalog cache was referenced to obtain table descriptor information. The value format is an integer.
- The type is int.
- The unit is lookups.

#### Application Locks Held

- The number of locks that are currently held. The value format is an integer.
- The type is int.
- The unit is locks.

#### Application Pool Read Time

- The total amount of elapsed time spent processing read requests that caused data or index pages to be physically read from disk to buffer pool. The value format is an integer.
- The type is int.
- The unit is milliseconds.

#### Application Failed SQL Stmts

- The number of SQL statements that were attempted, but failed. The value format is an integer.
- The type is int.
- The unit is statements.

#### Application Int Deadlock Rollbacks

- The total number of forced rollbacks initiated by the database manager due to a deadlock. The value format is an integer. The database manager initiates a rollback for the current unit of work

in an application that is experiencing a deadlock. This attribute shows the number of deadlocks that have been broken. It can indicate the possibility of concurrency problems. It is also important because internal rollbacks due to deadlocks can cause performance degradation.

- The type is int.
- The unit is rollbacks.

#### Application Direct Write Time

- The elapsed time (in milliseconds) required to perform the direct writes. The value format is an integer.
- The type is int.
- The unit is milliseconds.

#### Application Total Sorts

- The total number of sorts that have been issued. The value format is an integer. at the database or application level, use this value with the Sort Overflows attribute to calculate the percentage of sorts that need more heap space. You can also use it with the Total Sort Time attribute to calculate the average sort time. If the number of sort overflows is small with respect to the total sorts, increasing the sort heap size might have little impact on performance, unless this buffer size is increased substantially.
- The type is int.
- The unit is sorts.

#### Application Deadlocks

- The total number of deadlocks that have occurred. The value format is an integer. This attribute can indicate that applications are experiencing contention problems. To resolve the problem, determine in which applications (or application processes) the deadlocks are occurring. You can then modify the application to enable it to run concurrently. Some applications, however, might not be capable of running concurrently.
- The type is int.
- The unit is deadlocks.

#### Application Pool Total Writes

- The total number of write requests. The value format is an integer. This attribute is the total of the Pool Data Writes and Pool Index Writes attributes. Use this attribute to determine how busy the DB2 server is in terms of write I/O activity. Values that are greater than or equal to 2147483647 are indicated in the portal with the Value Exceeds Maximum text, and values that are smaller than -2147483648 are indicated with the Value Exceeds Minimum text.
- The type is int.
- The unit is requests.

#### Application Rej Curs Blk

- The number of times that a request for an I/O block at the server was rejected and the request was converted to non-blocked I/O. If there are many cursors blocking data, the communication heap might become full. The value format is an integer. When this heap is full, I/O blocks are not allocated for blocking cursors; however, an error condition does not alert you to this condition. If cursors are unable to block data, performance can be affected adversely.
- The type is int.
- The unit is occurrences.

#### Application Int Rows Deleted

- The number of rows deleted from the database as a result of internal activity. The value format is an integer. This attribute can help to gain insight into internal activity within the database manager. If this activity is high, you must evaluate the table design to determine if the referential constraints or triggers that you defined on the database are necessary.
- The type is int.
- The unit is rows.

#### Lock Conflict00 Lock Wait Time

- The total elapsed time (in milliseconds) that a lock was waited for. At the database level, this is the total amount of elapsed time that all applications were waiting for a lock within this database. At the application-connection and transaction levels, this is the total amount of elapsed time that this connection or transaction has waited for a lock to be granted. This attribute might be used with the Lock Waits attribute to calculate the average wait time for a lock. This calculation can be performed at either the database or the application-connection level. The following value is also valid:
- The type is int.
- The unit is milliseconds.

#### Application Pool Index to Estore

- Number of buffer pool index pages copied to extended storage. The value format is an integer. Pages are copied from the buffer pool to extended storage when they are selected as victim pages. As a result of the copying process, there is sufficient space for new pages in the buffer pool.
- The type is int.
- The unit is pages.

#### Application Pkg Cache Inserts

- The total number of times that a requested section was not available for use and had to be loaded into the package cache. The value format is an integer. This count includes any implicit prepares performed by the system.
- The type is int.
- The unit is inserts.

#### Application Pool Total Reads

- The total number of read requests that required I/O to get data pages and index pages into the buffer pool. The value format is an integer. This attribute is the total of the Pool Data Physical Reads and Pool Index Physical Reads attributes. Use this attribute to determine how busy the DB2 server is in terms of I/O activity. Values that are greater than or equal to 2147483647 are indicated in the portal with the Value Exceeds Maximum text, and values that are smaller than -2147483648 are indicated with the Value Exceeds Minimum text.
- The type is int.
- The unit is requests.

#### Application Appl Idle Time

- The number of seconds since an application issued a request to the server. The value format is an integer.
- The type is int.
- The unit is seconds.

#### Application Direct Read Reqs

- The number of requests to perform a direct read of one or more sectors of data. The value format is an integer.
- The type is int.
- The unit is requests.

#### Application Avg Lock Wait Time

- The average elapsed time (in milliseconds) that was spent waiting for a lock. The value format is an integer. If the average lock wait time is high, you must look for applications that hold many locks, or have lock escalations, with a focus on tuning your applications to improve concurrency, if appropriate. If escalations are the reason for a high average lock wait time, the values of one or both of the LOCKLIST and MAXLOCKS configuration parameters might be too low.
- The type is int.
- The unit is milliseconds.

#### Application Lock Timeouts

- The number of times that a request to lock an object time out instead of being granted. The value format is an integer.
- The type is int.
- The unit is timeouts.

#### Application Group00 Prefetch Wait Time

- The time an application spent waiting for an I/O server or prefetcher to finish loading pages into the buffer pool. This attribute can be used to experiment with changing the number of I/O servers and the I/O server sizes.
- The type is int.
- The unit is milliseconds.

#### Application Pool Index P Reads

- The number of physical read requests to get index pages into the buffer pool. The value format is an integer.
- The type is int.
- The unit is requests.

#### Application Pool Data L Reads

- The number of logical read requests for data pages that have gone through the buffer pool. The value format is an integer. This count includes accesses to data that is already in the buffer pool when the database manager needs to process the page or read into the buffer pool before the database manager can process the page.
- The type is int.
- The unit is requests.

#### Application Query Cost Estimate

- Estimated cost, in timerons, for a query, as determined by the SQL compiler. The value format is an integer. This attribute allows correlation of actual runtime values with the compile-time estimates.
- The type is int.
- The unit is timerons.

#### Application Int Rows Updated

- The number of rows updated from the database as a result of internal activity. The value format is an integer. This attribute can help to gain insight into internal activity within the database manager. If this activity is high, you must evaluate the table design to determine if the referential constraints that you defined on the database are necessary.
- The type is int.
- The unit is rows.

#### Application Rows Inserted

- The number of row insertions attempted. The value format is an integer. Use this attribute to gain insight into the current level of activity within the database manager.
- The type is int.
- The unit is inserts.

#### Application Rollback SQL Stmts

- The total number of SQL ROLLBACK statements that have been attempted. The value format is an integer. A rollback can result from an application request, a deadlock, or an error situation. This attribute counts only the number of rollback statements issued from applications.
- The type is int.
- The unit is rollbacks.

#### Application Total Hash Loops

- The total number of times that a single partition of a hash join was larger than the available sort heap space. The value format is an integer. Values for this attribute indicate inefficient execution of hash joins. This might indicate that the sort heap size is too small or the sort heap threshold is too small.
- The type is int.
- The unit is occurrences.

#### Application Open Rem Curs Blk

- The number of remote blocking cursors currently open for this application. The value format is an integer. Use this attribute with the Open Remote Cursors attribute to calculate the percentage of remote cursors that are blocking cursors.
- The type is int.
- The unit is cursors.

#### Application Pool Write Time

- The total amount of time spent physically writing data or index pages from the buffer pool to disk. The value format is an integer. Use this attribute with the Buffer Pool Data Writes and Buffer Pool Index Writes attributes to calculate the average page-write time. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device.
- The type is int.
- The unit is milliseconds.

#### Application Pool Index from Estore

- Number of buffer pool index pages copied from extended storage. The value format is an integer.
- The type is int.
- The unit is pages.

#### Application Internal Commits

- The total number of commits initiated internally by the database manager. The value format is an integer.
- The type is int.
- The unit is commits.

#### Application UID SQL Stmts

- The number of SQL UPDATE, INSERT, and DELETE statements that were issued. The value format is an integer.
- The type is int.
- The unit is statements.

#### Application Direct Reads

- The number of read operations that do not use the buffer pool. The value format is an integer.
- The type is int.
- The unit is reads.

#### Application Open Rem Curs

- The number of remote cursors currently open for this application, including the cursors counted by the Open Remote Cursors with Blocking attribute. The value format is an integer.
- The type is int.
- The unit is cursors.

#### Application Rows Updated

- The number of row updates attempted. The value format is an integer. Use this attribute to gain insight into the current level of activity within the database manager.
- The type is int.
- The unit is updates.

#### Application Query Card Estimate

- An estimate of the number of rows that are returned by a query. The value format is an integer. You can compare this estimate by the SQL compiler with the actual runtime values.
- The type is int.
- The unit is rows.

#### Application Dynamic SQL Stmts

- The number of dynamic SQL statements that were attempted. The value format is an integer.
- The type is int.
- The unit is statements.

#### Application Acc Curs Blk

- The number of times that a request for an I/O block was accepted. The value format is an integer. Use this attribute with the Rejected Block Cursor Requests attribute to calculate the percentage of blocking requests that are accepted or rejected.
- The type is int.
- The unit is accepteds.

#### Application Pool Data Writes



- The number of times a buffer pool data page was physically written to disk. The value format is an integer.
- The type is int.
- The unit is writes.

#### Application Group00 UOW Log Space Used

- The amount of log space (in bytes) used in the current unit of work of the monitored application. Use this attribute to understand the logging requirements at the unit-of-work level. Values that are greater than or equal to 2147483647 are indicated in the portal with the Value Exceeds Maximum text, and values that are smaller than -2147483648 are indicated with the Value Exceeds Minimum text.
- The type is int.
- The unit is bytes.

#### Application Failed SQL Stmts Percent

- The percentage of SQL statements that failed to run successfully. The value format is an integer. This value is derived by dividing the value of the Failed SQL Statements attribute by the value of the Total SQL Statements attribute.
- The type is double.
- The unit is percent.

#### Application Lock Waits

- The total number of times the database applications waited for locks. The value format is an integer. At the database level, the lock waits value is the total number of times that applications waited for locks within this database. At the application-connection level, the lock waits value is the total number of times that this connection requested a lock but waited because another connection was already holding a lock on the data.
- The type is int.
- The unit is occurrences.

#### Application Total Hash Joins

- The total number of hash joins that ran. The value format is an integer.
- The type is int.
- The unit is joins.

#### Application Direct Write Reqs

- The number of requests to perform a direct write of one or more sectors of data. The value format is an integer.
- The type is int.
- The unit is requests.

#### Application Lock Wait Time

- The total elapsed time (in milliseconds) that was spent waiting for a lock.
- The type is int.
- The unit is milliseconds.

#### Application Total SQL Stmt

- The total number of dynamic and static SQL statements. This value is derived by adding the values of the Dynamic SQL Statements and the Static SQL Statements attributes.

- The type is int.
- The unit is statements.

#### Application Sort Overflows Percent

- The percentage of sorts that ran out of sort heap space and might have required disk space for temporary storage. The value format is an integer. This percentage is calculated by dividing the value of the Sort Overflows attribute by the value of the Total Sorts attribute. at the database or application level, use this attribute to evaluate the percentage of sorts that required overflow to disk. If this percentage is high, you might want to adjust the database configuration by increasing the value of the SORTHEAP configuration parameter.
- The type is double.
- The unit is percent.

#### Application Degree Parallelism

- The degree of parallelism requested when the query was bound. The value format is an integer. Use with the Agents Top attribute to determine if the query achieved maximum level of parallelism.
- The type is int.
- The unit is degree.

#### Application Total Sort Time

- The total elapsed time (in milliseconds) for all sorts that ran. The value format is an integer. at the database or application level, use this element with the Total Sorts attribute to calculate the average sort time. This average can indicate whether sorting is a performance concern.
- The type is int.
- The unit is milliseconds.

#### Application DDL SQL Stmts

- The number of SQL Data Definition Language (DDL) statements that were issued. The value format is an integer.
- The type is int.
- The unit is statements.

### **Component: DB2 Table**

The Table data set provides information to monitor table-specific attributes, such as row read and row write rates.

#### **Dimensions**

##### Table Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string.

##### Table Schema

- The schema of the table for which the information is collected. The value format is a text string with a maximum of 60 bytes.
- The type is string.

##### Table DB Name U

- The database name. Data collection performance can be improved with the use of an eventing threshold or a filter in the Attribute Details tab that uses this attribute. When you specify a filter

that has a distinct database name or list of database names, the agent data collector filters the return data before the data is transmitted to the Performance Management console.

- The type is string. This is a key dimension.

#### Table DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. When specifying a filter with a distinct database partition value or list of values, the monitoring agent returns data for the requested partitions. If you do not specify a db partition filter value or if you specify a db partition filter value that is not valid, the return data is for the current partition. If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data collection includes both aggregated and individual partition attribute data.
- The type is string. This is a key dimension.

#### Table DB Name

- The database name. Data collection performance can be improved with the use of an eventing threshold or a filter in the Attribute Details tab that uses this attribute. When you specify a filter that has a distinct database name or list of database names, the agent data collector filters the return data before the data is transmitted to the Performance Management console.
- The type is string.

#### Table Name

- The name of the table for which the information is collected. The value format is a text string with a maximum of 60 bytes. Data collection performance can be improved with the use of an eventing threshold or a filter in the Attribute Details tab that uses this attribute. When you specify a filter that has a distinct database name or list of database names, the agent data collector filters the return data before the data is transmitted to the Performance Management console.
- The type is string. This is a key dimension.

#### Table Reorg Needed

- Indicates whether the table, its indexes, or both need to be reorganized, and is calculated using DB2 monitoring data that is generated when the DB2 RUNSTATS utility is run. The RUNSTATS utility collects statistics on tables and indexes, and can affect system performance as it is collecting the statistics. For this reason, RUNSTATS is not automatically run.
- The type is string.

#### Table Tablespace

- The name of the primary tablespace for the table. If no other tablespace is specified, all parts of the table are stored in this table space.
- The type is string.

#### Table Instance Name

- The name of the monitored DB2 instance.
- The type is string. This is a key dimension.

### Metrics

#### Table XML Object

- The disk space, in kilobytes, that is logically and physically allocated for the XML data in a table.
- The type is int.
- The unit is kilobytes.

#### Table Rows Write Rate for Interval

- The rate (per second) at which rows were changed (inserted, deleted, or updated) in the table during the monitoring interval.
- The type is double.
- The unit is rows/second.

#### Table Rows Read Rate for Interval

- The rate (per second) at which rows were read from the table during the monitoring interval.
- The type is double.
- The unit is rows/second.

#### Table Index Object Size

- The disk space, in kilobytes, that is logically and physically allocated for the indexes defined on the table.
- The type is int.
- The unit is kilobytes.

#### Table LOB Object

- The disk space, in kilobytes, that is logically and physically allocated for large objects in a table.
- The type is int.
- The unit is kilobytes.

#### Table Data Object Size

- The disk space, in kilobytes, that is logically and physically allocated for the table.
- The type is int.
- The unit is kilobytes.

### **Component: Database Activities**

Information about database activities.

#### **Dimensions**

##### Activities Instance Name

- The name of the monitored DB2 instance.
- The type is string. This is a key dimension.

##### Activities Snapshot Timestamp

- The date and time when the database system monitored information was collected. Use this attribute to help correlate data chronologically if you are saving the results in a file or database for ongoing analysis.
- The type is timestamp.

##### Activities DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If a db partition filter is not specified, data is returned for the current database partition. If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range.

- The type is string. This is a key dimension.

#### Activities Node Name

- The format is instanceid:hostname:UD for all operating systems. The format for version 6, release 1 of the DB2 agent Windows systems is instanceid:hostname:UD; on UNIX and Linux systems, the format is instanceid:hostname.
- The type is string.

#### Activities Catalog Partition

- The number of the catalog node.
- The type is int.

#### Activities Last Reset

- Indicates the most recent date and time when the monitor counters are reset for the application issuing the GET SNAPSHOT command. Use this attribute to determine the scope of information returned by the database system monitor.
- The type is timestamp.

#### Activities DB Name

- The real name of the host database for which information is collected or to which the application is connected. This name was given to the database when it was created. The value format is a simple text string with a maximum of 60 bytes.
- The type is string. This is a key dimension.

#### Activities Smallest Log Avail Node

- Indicates the node with the least amount (in bytes) of available log space.
- The type is int.

#### Activities Num DB Storage Paths

- The number of automatic storage paths that are associated with this database.
- The type is int.

#### Activities Appl ID Oldest Xact

- The application ID of the application that has the oldest transaction. Use this attribute to which application has the oldest active transaction. This application can be forced to free up log space. If the application take a large amount of log space, examine the application to determine if it can be modified to commit more frequently.
- The type is int.

#### Activities Catalog Partition Name

- The network name of the catalog node.
- The type is string.

#### Activities Sort Shrheap Allocated

- The total amount of shared sort memory allocated in the database.
- The type is int.

### Metrics

#### Activities Priv Workspace Section Lookups

- Indicates how many times the private workspace was accessed in order to locate a specific section for an application. At the database level, it is the cumulative total of all lookups for every application across all private workspaces in the database. At the application level, it is the cumulative total of all lookups for all sections in the private workspace for this application.
- The type is int.
- The unit is accesses.

#### Activities Priv Workspace Size Top

- The largest size reached by the private workspace.
- The type is int.
- The unit is bytes.

#### Activities Num Indoubt Trans

- The number of outstanding indoubt transactions in the database. Indoubt transactions hold log space for uncommitted transactions, which can cause the logs to become full. When the logs are full, further transactions can not be completed. The resolution of this problem involves a manual process of heuristically resolving the indoubt transactions. This attributes provides a count of the number of currently outstanding indoubt transactions that must be heuristically resolved.
- The type is int.
- The unit is transactions.

#### Activities Active Hash Joins

- The total number of hash joins that are currently running and consuming memory.
- The type is int.
- The unit is joins.

#### Activities Blocks Pending Cleanup

- The total number of MDC table blocks in the database that are pending asynchronous cleanup following a roll out delete.
- The type is int.
- The unit is tableBlocks.

#### Activities Stats Fabrications

- The total number of statistics fabrications that are performed by real-time statistics during query compilation for all the database applications. Instead of obtaining statistics by scanning data stored in a table or an index, statistics are fabricated based on metadata maintained by the index and data manager. Values reported by all the database partitions are aggregated together.
- The type is int.
- The unit is fabrications.

#### Activities Database Activities Log to Redo for Recovery

- The amount of log (in bytes) that has to be redone for crash recovery.
- The type is int.
- The unit is bytes.

#### Activities Pool Temp Index P Reads

- The number of index pages read in from the tablespace containers (physical) for temporary tablespaces.
- The type is int.

- The unit is pages.

#### Activities Cat Cache Size Top

- The largest size that is reached by the catalog cache. This attribute indicates the maximum number of bytes the catalog cache required for the workload run against the database since it is activated. If the catalog cache overflows, the value is the largest size reached by the catalog cache during the overflow.
- The type is int.
- The unit is bytes.

#### Activities Pool XDA P Reads

- The number of data pages for XML storage objects (XDAs) that are read in from the tablespace containers (physical) for regular and large tablespaces.
- The type is int.
- The unit is pages.

#### Activities Shr Workspace Size Top

- The largest size reached by shared workspaces.
- The type is int.
- The unit is bytes.

#### Activities Log Write Time S

- The total elapsed time (in seconds) that the logger spends writing log data to the disk. Use this attribute with the log writes and num log write io attributes to determine whether the current disk is adequate for logging.
- The type is int.
- The unit is seconds.

#### Activities Log Write Time NS

- The total elapsed time (in ns) that the logger spends writing log data to the disk. Use this attribute with the log writes and num log write io attributes to determine whether the current disk is adequate for logging.
- The type is int.
- The unit is nanoseconds.

#### Activities Database Activities Num Log Data Found in Buffer

- The number of times that an agent reads log data from the buffer. Reading log data from the buffer is preferable to reading from the disk because the latter is slower. Use this attribute with the num log read io attribute to determine if the LOGBUFSZ database configuration parameter needs to be increased.
- The type is int.
- The unit is occurrences.

#### Activities Stats Fabricate Time

- The total time (in milliseconds) spent on statistics fabrications by real-time statistics gathering. Statistics fabrication is the statistics collection activity needed to generate statistics during query compilation. If this attribute value is collected at the database level, it represents the total time spent on real-time statistics gathering activities for all the applications running on the database. If this attribute value is collected at the statement level, it represents the time spent

on the latest real-time statistics gathering activities for the statement. The times reported by all the database partitions are aggregated together.

- The type is int.
- The unit is milliseconds.

#### Activities Unread Prefetch Pages

- The number of pages that the prefetcher read in but are never used. If the value is high, prefetchers are causing unnecessary I/O by reading pages into the buffer pool that will not be used.
- The type is int.
- The unit is pages.

#### Activities Sync Runstats Time

- The total time spent on synchronous RUNSTATS activities triggered by real-time statistics gathering, in milliseconds. The synchronous RUNSTATS activities occur during query compilation. At the database level, this attribute value represents the total time spent on synchronous RUNSTATS activities for all the applications running on the database, triggered by real-time statistics gathering. At the statement level, this attribute value represents the time spent on the latest synchronous RUNSTATS activities for a particular statement, triggered by real-time statistics gathering. Values reported by all the database partitions are aggregated together.
- The type is int.
- The unit is milliseconds.

#### Activities Database Activities Num Log Part Page IO

- The number of I/O requests that are issued by the logger for writing partial log data to the disk. Use this attribute with the log writes, log write time, and num log write io attributes to determine if the current disk is adequate for logging. The following value is valid:
- The type is int.
- The unit is requests.

#### Activities Pool Temp Index L Reads

- The number of index pages that have been requested from the buffer pool (logical) for temporary tablespaces.
- The type is int.
- The unit is pages.

#### Activities Log Read Time NS

- The total elapsed time (in ns) that the logger spends reading log data from the disk. Use this attribute with the log reads, num log read io, and num log data found in buffer attributes to determine the following items:
- The type is int.
- The unit is nanoseconds.

#### Activities Pool Temp XDA L Reads

- The number of pages for XML storage object (XDA) Data that is requested from the buffer pool (logical) for temporary tablespaces.
- The type is int.
- The unit is pages.



#### Activities Shr Workspace Num Overflows

- The number of times that shared workspaces overflowed the bounds of their allocated memory. Use this attribute with the shr workspace size top attribute to determine whether the size of the shared workspaces need to be increased to avoid overflowing. Overflows of shared workspaces might cause performance degradation and out of memory errors from the other heaps that are allocated out of application shared memory.
- The type is int.
- The unit is overflows.

#### Activities Sync Runstats

- The total number of synchronous RUNSTATS activities triggered by real-time statistics gathering for all the applications in the database. This value includes both successful and unsuccessful synchronous RUNSTATS commands. Values reported by all the database partitions are aggregated together.
- The type is int.
- The unit is activities.

#### Activities Database Activities Num Log Write IO

- The number of I/O requests that are issued by the logger for writing log data to the disk. Use this attribute with the log writes and log write time attributes to determine if the current disk is adequate for logging.
- The type is int.
- The unit is requests.

#### Activities Priv Workspace Section Inserts

- The number of inserts of SQL sections by applications into the private workspace. The working copy of executable sections are stored in the private workspace. This attribute indicates the number of times when a copy was not available and had to be inserted. At the database level, it is the cumulative total of all inserts for every application across all private workspaces in the database. At the application level, it is the cumulative total of all inserts for all sections in the private workspace for this application.
- The type is int.
- The unit is inserts.

#### Activities Min Catalog Cache Size

- The minimum size of the catalog cache that is required by your workload.
- The type is int.
- The unit is pages.

#### Activities Post Shr Threshold Sorts

- The total number of sorts that were throttled back by the sort memory throttling algorithm. A throttled sort is a sort that was granted less memory than requested by the sort memory manager. A sort is throttled back when the memory allocation for sorts is close to the limit that is set by the sheapthres\_shr database configuration parameter. This throttling significantly reduces the number of overflows over sheapthres\_shr limit in a system that is not properly configured. The data reported by this attribute only reflects sorts using memory allocated from the shared sort heap.
- The type is int.
- The unit is sorts.

#### Activities Database Activities Num Log Buffer Full

- The number of times that agents have to wait for log data to write to disk while copying log records into the log buffer. This value is increased per agent per incident. For example, if two agents attempt to copy log data while the buffer is full, this value is increased by two. Use this attribute to determine if the LOGBUFSZ database configuration parameter needs to be increased.
- The type is int.
- The unit is occurrences.

#### Activities Database Activities Num Log Read IO

- The number of I/O requests that are issued by the logger for reading log data from the disk. Use this attribute with the log reads and log read time attributes to determine if the current disk is adequate for logging. The following value is valid:
- The type is int.
- The unit is requests.

#### Activities Pool Temp Data L Reads

- The number of data pages that have been requested from the buffer pool (logical) for temporary tablespaces. In conjunction with the pool temp data p reads attribute, the data page hit ratio for buffer pools located in temporary tablespaces can be calculated using the following formula:  $1 - (\text{pool temp data p reads} / \text{pool temp data l reads})$ .
- The type is int.
- The unit is pages.

#### Activities Priv Workspace Num Overflows

- The number of times that the private workspaces overflowed the bounds of its allocated memory. Use this attribute with the priv workspace size top attribute to determine whether the size of the private workspace needs to be increased to avoid overflowing. Overflows of the private workspace might cause performance degradation and out of memory errors from the other heaps allocated out of agent private memory.
- The type is int.
- The unit is overflows.

#### Activities Num Threshold Violations

- The number of threshold violations that have taken place in this database since the database was last activated. Use this attribute to determine whether thresholds are effective for this particular application or whether the threshold violations are excessive.
- The type is int.
- The unit is violations.

#### Activities Total OLAP Funcs

- The total number of OLAP functions that run.
- The type is int.
- The unit is functions.

#### Activities Elapsed Exec Time MS

- At the DCS statement level, this is the elapsed time (in ms) spent processing an SQL request on a host database server.
- The type is int.
- The unit is milliseconds.

#### Activities Stats Cache Size

- The current size of the statistics cache, which is used in a catalog partition to cache statistics information generated by real-time statistics gathering. Use this attribute to determine the size of the current statistics cache.
- The type is int.
- The unit is bytes.

#### Activities Shr Workspace Section Lookups

- Indicates how many times shared workspaces were accessed in order to locate a specific section for an application. At the database level, it is the cumulative total of all lookups for every application across all shared workspaces in the database. At the application level, it is the cumulative total of all lookups for all sections in the shared workspace for this application.
- The type is int.
- The unit is lookups.

#### Activities Database Activities Log Held By Dirty Pages

- The amount of log (in bytes) corresponding to the difference between the oldest dirty page in the database and the top of the active log. When the snapshot is taken, this value is calculated based on conditions at the time of that snapshot. Use this element to evaluate the effectiveness of page cleaning for older pages in the buffer pool.
- The type is int.
- The unit is bytes.

#### Activities Elapsed Exec Time S

- At the DCS statement level, this is the elapsed time (in seconds) spent processing an SQL request on a host database server.
- The type is int.
- The unit is seconds.

#### Activities Data Temp Pool Hit Ratio

- The data page hit ratio for buffer pools that are located in temporary tablespaces.
- The type is double.
- The unit is percent.

#### Activities OLAP Func Overflows

- The number of times that OLAP function data exceeded the available sort heap space. At the database level, use this attribute in conjunction with the total olapfuncs attribute to calculate the percentage of OLAP functions that overflowed to disk. If this percentage is high and the performance of applications using OLAP functions needs to be improved, consider increasing the sort heap size. At the application level, use this attribute to evaluate OLAP function performance for individual applications.
- The type is int.
- The unit is overflows.

#### Activities Pool No Victim Buffer

- The number of times an agent does not have a preselected victim buffer that is available. Use this attribute to help evaluate whether you have enough page cleaners for a given buffer pool when using proactive page cleaning.
- The type is int.

- The unit is occurrences.

#### Activities Pool Temp XDA P Reads

- The number of pages for XML storage object (XDA) Data that is read in from the tablespace containers (physical) for temporary tablespaces.
- The type is int.
- The unit is pages.

#### Activities Pool XDA L Reads

- The number of data pages for XML storage objects (XDAs) that are requested from the buffer pool (logical) for regular and large tablespaces.
- The type is int.
- The unit is pages.

#### Activities Post Shr Threshold Hash Joins

- The total number of hash joins that were throttled back by the sort memory throttling algorithm. A throttled hash join is a hash join that was granted less memory than requested by the sort memory manager. A hash join is throttled back when the memory allocation from the shared sort heap is close to the limit that is set by the sheapthres\_shr database configuration parameter. This throttling significantly reduces the number of overflows over sheapthres\_shr limit in a system that is not properly configured. The data reported in this element only reflects hash joins using memory allocated from the shared sort heap.
- The type is int.
- The unit is joins.

#### Activities Database Activities Total Log Used Pct

- The percentage of the log space that is used in the database.
- The type is double.
- The unit is percent.

#### Activities Pool XDA Writes

- The number of times that a buffer pool data page for an XML storage object (XDA) is physically written to disk.
- The type is int.
- The unit is occurrences.

#### Activities Pkg Cache Num Overflows

- The number of times that the package cache overflowed the bounds of its allocated memory.
- The type is int.
- The unit is occurrences.

#### Activities Async Runstats

- The total number of successful asynchronous RUNSTATS activities that are performed by real-time statistics gathering for all the applications in the database. Values reported by all the database partitions are aggregated together.
- The type is int.
- The unit is activities.

#### Activities database Activities Total Log Available

- The amount of active log space in the database that is not being used by uncommitted transactions (in bytes). Use this element in conjunction with the total log used attribute to determine whether you need to adjust the following configuration parameters of the monitored DB2 instance to avoid running out of log space:
- The type is int.
- The unit is bytes.

#### Activities Pool Temp Data P Reads

- The number of data pages read in from the tablespace containers (physical) for temporary tablespaces.
- The type is int.
- The unit is pages.

#### Activities Pkg Cache Size Top

- The largest size that is reached by the package cache. If the package cache overflowed, this attribute value is the largest size that is reached by the package cache during the overflow. Check the pkg cache num overflows attribute to determine if such a condition occurred.
- The type is int.
- The unit is bytes.

#### Activities Rows Read

- The number of rows that are read from the table. Use this attribute to identify tables with heavy usage, and for which you might want to create additional indexes. This attribute is not the number of rows that are returned to the calling application; it is the number of rows that must be read in order to return the result set.
- The type is int.
- The unit is rows.

#### Activities Log Read Time S

- The total elapsed time (in seconds) that the logger spends reading log data from the disk. Use this attribute with the log reads, num log read io, and num log data found in buffer attributes to determine the following items:
- The type is int.
- The unit is seconds.

#### Activities Active OLAP Funcs

- The total number of OLAP functions that are currently running and consuming sort heap memory.
- The type is int.
- The unit is functions.

#### Activities Min Pkg Cache Size

- The minimum size of the package cache that is required by your workload.
- The type is int.
- The unit is pages.

#### Activities Sort Shrheap Top

- The high watermark (in 4KB pages) of the database-wide shared sort memory.
- The type is int.

- The unit is 4kilobytePages.

#### Activities Shr Workspace Section Inserts

- The number of inserts of SQL sections by applications into shared workspaces. The working copy of executable sections are stored in shared workspaces. This attribute indicates the number of times when a copy was not available and had to be inserted. At the database level, it is the cumulative total of all inserts for every application across all shared workspaces in the database. At the application level, it is the cumulative total of all inserts for all sections in the shared workspace for this application.
- The type is int.
- The unit is inserts.

#### Activities Pool Temp Hit Ratio

- The data page and index page hit ratio for buffer pools that are located in temporary tablespaces.
- The type is double.
- The unit is percent.

### **Component: Events**

Detailed information about predefined and triggered events.

#### **Dimensions**

##### Event Subcategory

- The subcategory of the event.
- The type is string.

##### Event Suggestion

- The suggestion, detailing how best to proceed once the event has been triggered.
- The type is string.

##### Event Error Message

- The error message returned by the DB2 instance.
- The type is string.

##### Event DB Name

- The real name of the database for which information is collected. This name was given to the database when it was created. The value format is a simple text string with a maximum of 60 bytes. Use this attribute to identify the specific database to which the data applies.
- The type is string. This is a key dimension.

##### Event Host Name

- The hostname of the machine where the DB2 database is hosted.
- The type is string.

##### Event Time Stamp

- The local time on the agent when the event was triggered.
- The type is timestamp.

##### Event Level

- The level of the event. This can be either Error, Warning, Info or Misc.
- The type is int.

#### Event Node Name

- The managed system name of the agent. For new installations of version 7, release 1, the format is instanceid:hostname:UD for all operating systems.
- The type is string.

#### Event SQL State

- The SQL state returned by the DB2 instance.
- The type is string.

#### Event Error Code

- The error code returned by the DB2 instance.
- The type is int.

#### Event Instance Name

- The name of the monitored DB2 instance.
- The type is string.

#### Event Category

- The category of the event.
- The type is string.

#### Event Description

- The description of the event.
- The type is string.

### **DB2 Instance**

Information about the DB2 Instance.

#### **Dimensions**

##### Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string. This is a key dimension.

##### Product Version

- The product and version that is running on the DB2 instance.
- The type is string.

##### Req IO Blk

- The current value (in byte units) of the client input and output block size. This value is the amount of memory that is allocated for the communication buffer between remote applications and their database agents on the database server. When a database client requests a connection to a remote database, this communication buffer is allocated on the client. On the database server, a communication buffer of 32767 bytes is initially allocated, until a connection is established and the server can determine the value of the rqrioblk attribute at the client. In addition to this communication buffer, this parameter is also used to determine the input and output block size at the database client when a blocking cursor is opened.

- The type is int.

#### DB2 Start Timestamp

- The date and time that the database manager was started using the DB2START command. Use this attribute with the snapshot time attribute to calculate the elapsed time from the start of the database manager until the snapshot was taken.
- The type is timestamp.

#### Instance Name

- The name of the monitored DB2 instance.
- The type is string. This is a key dimension.

#### Max Conc Agents

- The maximum number of database manager coordinator agents that can concurrently run a database manager transaction in the DB2 instance during the monitoring interval. When this monitor is used with DB2 Universal database servers, the maxcagents value is the default. This value is the maximum number of database manager agents that can be concurrently executing a database manager transaction. Use the maxcagents attributes to control the load on the system during periods of high simultaneous application activity. A value of -1 indicates that the limit is equal to the maximum number of agents (the MAXAGENTS parameter). The maxcagents parameter does not limit the number of applications that can have connections to the database.
- The type is int.

#### Snapshot Timestamp

- The date and time when the database system monitored information was collected. Use this attribute to help correlate data chronologically if you are saving the results in a file or database for ongoing analysis.
- The type is timestamp.

#### DB2 Version

- The version of the server that is returning the data. For example: 6. 1 or 7. 1. The data structures used by the monitor might change between releases. As a result, check the version of the data stream to determine whether your applications can process the data.
- The type is string.

#### Query Heap Size

- The maximum amount of memory that can be allocated for the query heap within the DB2 instance during the monitoring interval. Use a query heap to store each query in the private memory of the agent. Use the results from the aslheapsz attribute to refine the query heap size.
- The type is int.

#### Max Agents

- The current value of the maximum number of existing agents. This value is the maximum number of database manager agents available at any given time to accept application requests. This value limits the total number of applications that can connect to all databases in the DB2 instance at a given time. The value of the maxagents attribute must be the sum of the values of the maxappls attribute in each database that is allowed to be accessed concurrently. Increasing the value of the maxagents attribute can increase resource use because resources for each agent are allocated when the DB2 instance is started. The following value is valid:
- The type is int.

#### Max Coord Agents



- The maximum number of database manager coordinating agents that can exist on a server in a partitioned or nonpartitioned database environment. One coordinating agent is acquired for each local or remote application that connects to a database or attaches to an instance. Requests that require an instance attachment include CREATE DATABASE, DROP DATABASE, and Database System Monitor commands.
- The type is int.

#### Mon Heap Size

- The current value (in units of 4-KB pages) of the database system monitor heap size. This value is the amount of memory that is allocated for database system monitor data. A value of zero prevents the database manager from collecting database system monitor data.
- The type is int.

#### DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If you do not specify a db partition filter, data is returned for either the current database partition (single partition environment) or the aggregated database partitions (multiple partition environment). If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range, the following values are also valid:
- The type is string. This is a key dimension.

#### Priority of Agents

- The current value of the priority of agents. This value is the priority that the operating system scheduler gives to agent and other database manager instance processes and threads. This priority determines how the operating system gives CPU time to the DB2 processes and threads relative to the other processes and threads running on the system. A value of -1 indicates that no special action is taken and the operating system schedules the database manager in the normal way that it schedules all processes and threads. Any other value indicates that the database manager creates its processes and threads with a static priority set to this value.
- The type is int.

#### Connection Status

- The status of the communication connection between the database partition that is specified by the DB2 Node Number variable and the database partition where this monitor runs. Two nodes can be active, but the connection between them remains inactive unless there is active communication between them.
- The type is int.

#### DBPG Node Status

- The list of failing local nodes. This is a list of integers, where each integer represents the failed local nodes. Depending on the actual partitions defined in the database partition group, it is not necessarily true that all the nodes defined in the parallel environment are examined. To ensure that all the nodes in the partitioned environment are examined, define a partition group that contains at least one database partition from each of the nodes in the partitioned environment.
- The type is string.

#### Last Reset Timestamp

- The date and time that the monitor counters were reset for the application requesting the snapshot. Use this attribute to help you determine the scope of information returned by the database system monitor.

- The type is timestamp.

#### Sort Heap Thres

- The current value (in units of 4-KB pages) of the sort heap threshold. This value is the maximum amount of memory that the database manager allocates for piped sorts. Piped sorts perform better than non-piped sorts and are used more often. However, their use can affect the performance. The value of the sheaphres attribute must be at least two times the largest sort heap that is defined for any database within the instance. The following value is valid:
- The type is int.

#### DB2 Instance Status

- The current status of the DB2 instance. Use this attribute to determine the state of your database manager instance.
- The type is string.

#### DB2 Server Type

- The type of database manager being monitored.
- The type is string.

### Metrics

#### Gateway Current Connections

- The current number of connections to host databases being handled by the DB2 Connect gateway. Use this attribute to help you understand the level of activity at the DB2 Connect gateway and the associated use of system resources. The following value is valid:
- The type is int.
- The unit is connections.

#### Remote Connections Executing

- The number of remote applications currently connected to a database and currently processing a unit of work within the database manager instance being monitored. By using this number, you can determine the level of concurrent processing occurring on the database manager. This value changes frequently. As a result, you must sample the data at specific intervals over an extended period of time to get a realistic view of system usage. This number does not include applications that were initiated from the same instance as the database manager. The following value is valid:
- The type is int.
- The unit is connections.

#### Max Agent Overflows

- The number of attempts to create a new agent when the MAXAGENTS configuration parameter had already been reached. If requests to create new agents are received after reaching the MAXAGENTS configuration parameter, the workload for this node might be too high. The following value is valid:
- The type is int.
- The unit is attempts.

#### Agents Waiting on Token Pct

- The percentage of agents waiting on a token. The percentage is calculated by dividing the value of the Agents Waiting on Token attribute by the number of local applications that are currently connected to a database (Local Cons attribute). Use this attribute to assess the number of agents.
- The type is double.
- The unit is percent.

### Remote Connections

- The current number of connections initiated from remote clients to the instance of the database manager that is being monitored. This attribute shows the number of connections from remote clients to databases in this instance. This value changes frequently. As a result, you must sample the data at specific intervals over an extended period of time to get a realistic view of system usage. This number does not include applications that were initiated from the same instance as the database manager. The following value is valid:
- The type is int.
- The unit is connections.

### RB Used Percent

- The percentage of FCM request blocks used within the partitioned database server during the monitoring interval. If the percentage of FCM request blocks used is high compared to normal operating levels, you can adjust the `fc_num_rqb` attribute.
- The type is double.
- The unit is percent.

### CE Max Used Percent

- The maximum percentage of FCM connection entries used during processing within the partitioned database server. If the percentage of maximum FCM connection entries used is high compared to normal operating levels, you can increase the number of FCM connections; if the percentage is low compared to normal operating levels, you can decrease the value.
- The type is double.
- The unit is percent.

### FCM Num Rqb

- The number of FCM request blocks for the DB2 instance during the monitoring interval. Request blocks are the media through which information is passed between the FCM daemon and an agent. The requirement for request blocks varies according to the number of users on the system, the number of database partition servers in the system, and the complexity of queries that are run. The following value is valid:
- The type is int.
- The unit is requestBlocks.

### Agents Registered Top

- The maximum number of agents that the database manager has ever registered, at the same time, since it was started (coordinator agents and subagents). Use this attribute to evaluate the setting of the `MAXAGENTS` configuration parameter. The number of agents registered at the time the snapshot was taken is recorded by the `Agents Registered` attribute. The following value is valid:
- The type is int.
- The unit is agents.

### Gateway Cons Wait Host

- For host databases being handled by the DB2 Connect gateway, the current number of connections that are waiting for a reply from the host. Because this value can change frequently, take samples at regular intervals over an extended period to obtain a realistic view of gateway usage. The following value is valid:
- The type is int.
- The unit is connections.

#### MA Free Bottom

- The minimum number of free message anchors. The following value is valid:
- The type is int.
- The unit is anchors.

#### Post Threshold Sorts

- The number of sorts that have requested heaps after reaching the sort heap threshold. By modifying the sort heap threshold and sort heap size configuration parameters, you can improve the performance of sort operations or the overall system. If the value of this attribute is high, you can do one of the following actions:

- Increase the sort heap threshold (sheapthres).
- Adjust applications to use fewer or smaller sorts by using SQL query changes.

The following value is valid:

- The type is int.
- The unit is sorts.

#### RB Max Used Percent

- The percentage of maximum FCM request blocks used during processing within the partitioned database server. If the percentage of maximum FCM request blocks used is high compared to normal operating levels, you can adjust the fcm\_num\_rqb attribute.
- The type is double.
- The unit is percent.

#### CE Used Percent

- The percentage of FCM connection entries used during processing within the partitioned database server. If the percentage of FCM connection entries used is high compared to normal operating levels, you can increase the number of FCM connections; if the percentage is low compared to normal operating levels, you can decrease the value.
- The type is double.
- The unit is percent.

#### Local Connections

- The number of local applications that are currently connected to a database within the database manager instance being monitored. By using this number, you can determine the level of concurrent processing occurring in the database manager. This value changes frequently. As a result, you must sample the data at specific intervals over an extended period of time to get a realistic view of system usage. This number includes only applications that were initiated from the same instance as the database manager. The applications are connected, but might or might not be executing a unit of work in the database. The following value is valid:
- The type is int.
- The unit is connections.

#### Total Buffers Rcvd

- The total number of FCM buffers received by the database node where this monitor runs. The database node is specified in the DB2\_node\_number variable. Use the returned value to measure the level of traffic between the node where this monitor runs and another node. If the total number of FCM buffers received from the other node is high compared to normal operating levels, you can redistribute the database or move tables to reduce the internode traffic. The following value is valid:
- The type is int.

- The unit is buffers.

#### FCM Num Buffers

- The number of buffers that are used for internal communications (messages) among the nodes and within the nodes in a DB2 instance during the monitoring interval. You might need to increase the value of this parameter if you have either of the following conditions: multiple logical nodes on a processor, or too many users, nodes, or complex applications that exceed the buffer limit. The following value is valid:
  - The type is int.
  - The unit is buffers.

#### CE Free Bottom

- The minimum number of free connection entries. The following value is valid:
  - The type is int.
  - The unit is connections.

#### Appl Support Layer Heap Size

- The current value (in units of 4-KB pages) of the application support layer heap size. This value is the amount of memory that is allocated for the application support layer heap. This heap is used as a communication buffer between the local application and its associated agent. In addition, this value is used to determine the input and output block size when a blocking cursor is opened. The following value is valid:
  - The type is int.
  - The unit is occurrences.

#### Total Buffers Sent

- The total number of FCM buffers that are sent from the database node where this monitor runs to the specified node. Use the returned value to measure the level of traffic between the current node where this monitor runs and the specified node. If the total number of FCM buffers sent to the other node is high compared to normal operating levels, you can redistribute the database or move tables to reduce the internode traffic. The following value is valid:
  - The type is int.
  - The unit is buffers.

#### Agents Created Empty Pool Ratio

- The percentage of agents that are created because the pool is empty. This ratio is calculated by dividing the value of the Agents Created Empty Pool attribute by the value of the Agents From Pool attribute. Use this attribute to evaluate how often an agent must be created because the pool is empty.
  - The type is double.
  - The unit is percent.

#### Gateway Total Connections

- The total number of connections attempted from the DB2 Connect gateway since the last db2start command or the last reset. Use this attribute to help you understand the level of activity at the DB2 Connect gateway and the associated use of system resources. The following value is valid:
  - The type is int.
  - The unit is connections.

#### Cons in Exec Percent

- The percentage of the maximum number of applications allowed that are connected to a database and processing a unit of work during the monitoring interval.
- The type is double.
- The unit is percent.

#### RB Free

- The number of request blocks that are free in the partitioned database server during the monitoring interval. Use the returned value with the `fcm_num_rqb` attribute to determine the current request block utilization. You can use this information to refine the `fcm_num_rqb` attribute. The following value is valid:
- The type is int.
- The unit is requestBlocks.

#### Piped Sorts Accepted Percent

- The percentage of piped sorts that have been accepted. The percentage is calculated by dividing the value of the Piped Sorts Accepted attribute by the value of the Piped Sorts Requested attribute. Use this attribute to determine whether the value of the Piped Sorts Accepted attribute is in an acceptable range.
- The type is double.
- The unit is percent.

#### Local Connection Executing

- The number of local applications that are currently connected to a database within the database manager instance being monitored and are currently processing a unit of work. By using this number, you can determine the level of concurrent processing occurring in the database manager. This value changes frequently. As a result, you must sample the data at specific intervals over an extended period of time to get a realistic view of system usage. This number includes only applications that were initiated from the same instance as the database manager. The following value is valid:
- The type is int.
- The unit is applications.

#### Agents Waiting on Token

- The number of agents waiting for a token so they can run a transaction in the database manager. Use this attribute to evaluate your setting for the `MAXCAGENTS` configuration parameter. Each application has a dedicated coordinator agent to process database requests (transactions) within the database manager. Each agent must have a token to run a transaction. The maximum number of coordinator agents is limited by the `MAXCAGENTS` configuration parameter. The following value is valid:
- The type is int.
- The unit is agents.

#### Agents Stolen

- The number of times that agents are stolen from an application. Agents are stolen when an idle agent associated with an application is reassigned to work on a different application. Use this attribute with the Maximum Number of Associated Agents attribute to evaluate the load that this application places on the system. The following value is valid:
- The type is int.
- The unit is agents.

#### Buff Free Bottom

- The minimum number of free connection entries. The following value is valid:
- The type is int.
- The unit is connections.

#### Sort Heap Allocated

- The total number of pages that are used for sorts at the selected level (database manager or database) when the snapshot was taken. Add the extra memory used for the sort heap to the base memory requirements for running the database manager when excessive sorting occurs. The large heap size and indexes can improve the sorting performance. The following value is valid:
- The type is int.
- The unit is pages.

#### Buff Max Used Percent

- The percentage of maximum FCM buffers used during processing within the partitioned database server. If the percentage of maximum FCM buffers used is high compared to normal operating levels, you can increase the number of FCM buffers; if the percentage is low compared to normal operating levels, you can decrease the value.
- The type is double.
- The unit is percent.

#### Agents Created Empty Pool

- The number of agents created because the agent pool was empty. It includes the number of agents started at DB2 start up. By using the Agents Assigned From Pool attribute, you can calculate the ratio of the Agents Created Empty Pool attribute to the Agents From Pool attribute. See the Agents From Pool attribute for information about using this attribute.
- The type is int.
- The unit is agents.

#### Agents Registered

- The number of agents that the database manager registered. The following value is valid:
- The type is int.
- The unit is agents.

#### FCM Num Connect

- The number of FCM connection entries for the DB2 instance during the monitoring interval. Agents use connection entries to pass data among themselves. Use the results from the fcm\_num\_rqb attribute to help you refine the fcm\_num\_connect attribute. The following value is valid:
- The type is int.
- The unit is connections.

#### Piped Sorts Accepted

- The number of piped sorts that have been accepted. When the number of accepted piped sorts is low compared to the number requested, you can improve sort performance by adjusting one or both of the following configuration parameters:
  - SORTHEAP
  - SHEAPTHRES

If piped sorts are being rejected, consider decreasing your sort heap or increasing your sort heap threshold. Be aware of the possible implications of these options:

- If you increase the sort heap threshold, more memory might remain allocated for sorting. This can cause the paging of memory to disk.
- If you decrease the sort heap, an extra merge phase (which can slow down the sort) might be required.

The following value is valid:

- The type is int.
- The unit is sorts.

#### Coordinating Agents Top

- The maximum number of coordinating agents working at one time. The MAXCAGENTS configuration parameter determines the number of coordinating agents that can be executing concurrently. If the peak number of coordinating agents results in a workload that is too high for this node, you can reduce the MAXCAGENTS configuration parameter. The following value is valid:
- The type is int.
- The unit is agents.

#### CE Free

- The number of connection entries that are free in the partitioned database server during the monitoring interval. Use the returned value to help determine the current connection entry utilization. The following value is valid:
- The type is int.
- The unit is connections.

#### RB Free Bottom

- The minimum number of free request blocks. The following value is valid:
- The type is int.
- The unit is requestBlocks.

#### CPU Used Pct

- The percentage of CPU that is used on the system by the DB2 instance. DB2 returns this value as SMALLINT.
- The type is int.
- The unit is percent.

#### Total Memory Used

- The total database memory used by instance.
- The type is int.
- The unit is MB.

#### Piped Sorts Requested

- The number of piped sorts that have been requested. Because piped sorts might reduce disk I/O, allowing more piped sorts can improve the performance of sort operations and possibly the performance of the overall system. A piped sort is not accepted if the sort heap threshold is exceeded by allocating the requested sort heap. See the Piped Sorts Accepted attribute for more information if piped sorts are being rejected. The SQL EXPLAIN output shows whether the optimizer requested a piped sort. The following value is valid:
- The type is int.
- The unit is sorts.



#### Agents from Pool

- The number of agents assigned from the pool. Use this attribute with Agents Created Empty Pool attribute to determine how often an agent must be created because the pool is empty. The following value is valid:
- The type is int.
- The unit is agents.

#### Post Threshold Hash Joins

- The total number of times that a hash join heap request was limited due to the concurrent use of shared or private sort heap space. If this value is large (for example, greater than 5% of Hash Join Overflows), you must consider increasing the sort heap threshold. The following value is valid:
- The type is int.
- The unit is occurrences.

#### MA Max Used Percent

- The maximum number of message anchors used as a percentage.
- The type is double.
- The unit is percent.

#### Total Memory Allocated

- The total memory allocated to the DB2 instance.
- The type is int.
- The unit is MB.

#### Post Threshold OLAP Funcs

- The number of OLAP functions that have requested a sort heap after the sort heap threshold has been exceeded. If the value of this attribute is high, increase the sort heap threshold (sheaphthres). The following value is valid:
- The type is int.
- The unit is functions.

#### Piped Sorts Rejected Percent for Interval

- The percentage of piped sort requests that were rejected for the DB2 instance during the monitoring interval. In the sort return phase, if the sorted information can return directly through the sort heap, it is a piped sort. However, even if the optimizer requests a piped sort, this request is rejected at run time if the total amount of sort heap memory for all sorts on the database is close to exceeding the sheaphthres value. If this returned value is high compared to normal operating levels, consider decreasing your sort heap (using the sortheap configuration parameter) or increasing your sort heap threshold (using the sheaphthres configuration parameter). However, be aware of the implications of these options. If you increase the sort heap threshold, more memory can remain allocated for sorting, causing the paging of memory to disk. If you decrease the sort heap, you can require an extra merge phase that can slow down the sort.
- The type is double.
- The unit is percent.

#### Agents Waiting Top

- The highest number of agents waiting on a token, at the same time, since the database manager was started. Use this attribute to evaluate the setting of the MAXCAGENTS configuration parameter. In contrast, the Agents Waiting on Token attribute records the number of agents waiting for a token at the time the snapshot was taken. The following value is valid:

- The type is int.
- The unit is agents.

#### Gateway Cons Wait Client

- For host databases being handled by the DB2 Connect gateway, the current number of connections that are waiting for the client to send a request. Because this value can change frequently, take samples at regular intervals over an extended period to obtain a realistic view of gateway usage. The following value is valid:
- The type is int.
- The unit is connections.

#### Buff Free

- The number of Fast Communication Manager (FCM) buffers that are free in the partitioned database server during the monitoring interval. Use the returned value to determine the current buffer pool utilization. Use this information to refine the configuration of the number of FCM buffers. The following value is valid:
- The type is int.
- The unit is buffers.

#### Sort Heap Used Percent

- The percentage of the allocated sort heap that the DB2 instance used during the monitoring interval.
- The type is double.
- The unit is percent.

#### Committed Private Memory

- The amount of private memory that the instance of the database manager currently has committed at the time of the snapshot. Use this attribute to assess the MIN\_PRIV\_MEM configuration parameter to ensure that enough private memory is available. This attribute is returned for all platforms, but tuning can be accomplished only on platforms where DB2 uses threads (such as OS/2 and Windows NT systems). Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal. The following value is valid:
- The type is int.
- The unit is kilobytes.

#### Conn Local Database

- The number of local databases with current connections to the monitored DB2 instance. This value gives an indication of how many database information records to expect when gathering data at the database level. The applications can be running locally or remotely, and might or might not be executing a unit of work within the database manager. The following value is valid:
- The type is int.
- The unit is databases.

#### DB2 Available

- The amount of time (in seconds) the instance has been available since a DB2START command was issued. The value format is an integer. The value is derived through this formula: snapshot time - db2start time The following value is valid:
- The type is int.
- The unit is seconds.

#### FCM Num Anchors

- The number of FCM message anchors for the DB2 instance during the monitoring interval. Agents use the message anchors to send messages among themselves. The following value is valid:
- The type is int.
- The unit is anchors.

#### Idle Agents

- The number of agents in the agent pool that are currently unassigned to an application. Use this attribute to set the NUM\_POOLAGENTS configuration parameter. By having idle agents available to satisfy requests for agents, you can improve performance. The following value is valid:
- The type is int.
- The unit is agents.

#### Buff Used Percent

- The percentage of FCM buffers that are used within the partitioned database server during the monitoring interval. If the percentage of FCM buffers used is high compared to normal operating levels, you can adjust the number of FCM buffers.
- The type is double.
- The unit is percent.

#### Piped Sort Hit Ratio Percent for Interval

- The piped sort hit ratio (as a percentage) for the last monitoring interval. The piped sort hit ratio is the ratio of piped sorts accepted to piped sorts requested.
- The type is double.
- The unit is percent.

#### Piped Sorts Rejected for Interval

- The total number of piped sorts that were rejected during the monitoring interval. In the return phase of sorting, if the sorted information can return directly through the sort heap, it is a piped sort. However, even if the optimizer requests a piped sort, this request is rejected at run time if the total amount of sort heap memory for all sorts on the database is close to exceeding the sheapthres value. If this returned value is high compared to the total number of sorts requested, consider decreasing your sort heap (using the sortheap configuration parameter) or increasing your sort heap threshold (using the sheapthres configuration parameter). However, be aware of the implications of these options. If you increase the sort heap threshold, more memory can remain allocated for sorting, causing the paging of memory to disk. If you decrease the sort heap, you can require an extra merge phase that can slow down the sort. The following value is valid:
- The type is int.
- The unit is sorts.

### **Component: Buffer Pool**

Information about buffer pool activities.

#### **Dimensions**

##### Buffer Pool Input DB Alias

- The alias of the database provided when calling the snapshot function. The value format is a simple text string with a maximum of 60 characters. Use this attribute to help you identify the specific database to which the monitor data applies. It contains blanks unless you requested monitor information related to a specific database.
- The type is string.

#### Buffer Pool BP Name

- The name of the buffer pool. A new database has a default buffer pool (named IBMDEFAULTBP). The size of the default buffer pool is determined by the platform. Depending on your needs you might choose to create several buffer pools, each of a different size, for a single database. The CREATE, ALTER, and DROP BUFFERPOOL statements allow you to create, change, or remove a buffer pool.
- The type is string.

#### Buffer Pool DB Path

- The full path of the location where the database is stored on the monitored system. Use this attribute with the Database Name attribute to identify the specific database to which the data applies.
- The type is string.

#### Buffer Pool DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If a db partition filter is not specified, data is returned for the current database partition. If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range, the following values are also valid:
- The type is string. This is a key dimension.

#### Buffer Pool Node Name

- The format is instanceid:hostname:UD for all operating systems. The format for version 6, release 1 of theDB2 agent on Windows systems is instanceid:hostname:UD; on UNIX and Linux systems, the format is instanceid:hostname.
- The type is string.

#### Buffer Pool DB Name

- The real name of the database for which information is collected or to which the application is connected. This name was given to the database when it was created. The value format is a simple text string with a maximum of 60 characters. Use this attribute to identify the specific database to which the data applies.
- The type is string. This is a key dimension.

### Metrics

#### Buffer Pool Pool Sync Data Writes

- The total number of physical write requests that were performed synchronously (that is, physical data page writes that were performed by database manager agents). This value is derived by subtracting the value of the Pool Async Data Writes attribute from the value of the Pool Data Writes attribute. By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing.
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Data L Reads

- The number of logical read requests for data pages that have gone through the buffer pool. This count includes accesses to the following data:
  - Data that is already in the buffer pool when the database manager needs to process the page.

- Data that is read into the buffer pool before the database manager can process the page.

By using the Pool Data Physical Reads attribute, you can calculate the data page hit ratio for the buffer pool as follows:  $1 - (\text{buffer pool data physical reads} / \text{buffer pool data logical reads})$  By using the Pool Data Physical Reads, Pool Index Physical Reads, and Pool Index Logical Reads attributes, you can calculate the overall buffer pool hit ratio as follows:  $1 - ((\text{buffer pool data physical reads} + \text{buffer pool index physical reads}) / (\text{buffer pool data logical reads} + \text{buffer pool index logical reads}))$  Increasing buffer pool size generally improves the hit ratio until you reach a point of diminishing returns.

- The type is int.
- The unit is reads.

#### Buffer Pool Avg Sync Write Time

- The average elapsed time used to perform a synchronous write. This value is derived by dividing the value of the Pool Sync Write Time attribute by the value of the Pool Sync Write attribute. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device.
- The type is int.
- The unit is ?.

#### Buffer Pool Logical Read Per Min

- The number of logical read operations that are performed on the buffer pool per minute.
- The type is int.
- The unit is reads/minute.

#### Buffer Pool Direct Reads

- The number of read operations that do not use the buffer pool. Use the following formula to calculate the average number of sectors that are read by a direct read:  $\text{direct reads from database} / \text{direct read requests}$  When using system monitors to track I/O, this data attribute helps to distinguish database I/O from non-database I/O on the device.
- The type is int.
- The unit is reads.

#### Buffer Pool Pool Index Writes

- The number of times a buffer pool index page was physically written to disk. If a buffer pool index page is written to disk for a high percentage of Buffer Pool Index Physical Reads, performance might improve by increasing the number of buffer pool pages available for the database. If all applications are updating the database, increasing the size of the buffer pool might have minimal impact on performance; most pages contain updated data that must be written to disk.
- The type is int.
- The unit is writes.

#### Buffer Pool Pool Async Write Time

- The total elapsed time spent writing data or index pages from the buffer pool to disk by database manager page cleaners. Calculate the elapsed time spent writing pages synchronously by subtracting the value of the Pool Async Write Time attribute from the value of the Pool Physical Write Time attribute. You can also use this attribute to calculate the average asynchronous read time:
  1. Sum the value of the Pool Async Data Writes attribute and the value of the Pool Async Index Writes attribute.

2. Divide the value of the Pool Async Write Time attribute by the sum from step 1.

These calculations can be used to understand the I/O work being performed.

- The type is int.
- The unit is ?.

#### Buffer Pool Pool Hit Ratio

- The buffer pool hit ratio (as a percentage). The sum of the Pool Data Logical Reads and Pool Index Logical Reads attributes is divided by the value of the Pool Total Reads attribute to derive the pool hit ratio. This attribute can determine whether buffer pool assignment is efficient. If the pool hit ratio is low, increasing the number of buffer pool pages might improve performance.
- The type is double.
- The unit is percent.

#### Buffer Pool Pool Sync Write

- The total number of synchronous index writes. The value is derived by adding the values of the Pool Sync Data Writes attribute and Pool Sync Index Writes attribute.
- The type is int.
- The unit is writes.

#### Buffer Pool Avg Pool Read Time

- The average elapsed time for a read request. This value is derived by dividing the value of the Pool Read Time attribute by the value of the Pool Total Reads attribute. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device.
- The type is int.
- The unit is ?.

#### Buffer Pool Pool Index from Estore

- Number of buffer pool index pages copied from extended storage. Required index pages are copied from extended storage to the buffer pool. The copy process might incur the cost of connecting to the shared memory segment, but it saves the cost of a disk read.
- The type is int.
- The unit is pages.

#### Buffer Pool Avg Direct Write Time

- The average elapsed time for a direct write request. This value is calculated by dividing the value of the Direct Write Time attribute by the value of the Direct Writes attribute. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device. The following value is also valid:
- The type is int.
- The unit is result.

#### Buffer Pool Async Write Ratio

- The ratio of buffer pool asynchronous data writes to the total number of pool writes for the database.
- The type is int.
- The unit is ratio.

#### Buffer Pool Direct Write Reqs

- The number of requests to perform a direct write of one or more sectors of data. Use the following formula to calculate the average number of sectors that are written by a direct write:  $\text{direct writes to database} / \text{direct write requests}$  The following value is also valid:
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Index L Reads

- The number of logical read requests for index pages that have gone through the buffer pool. This count includes accesses to the following index pages:
  - Pages that are already in the buffer pool when the database manager needs to process the page.
  - Pages that are read into the buffer pool before the database manager can process the page.

By using the Buffer Pool Index Physical Reads attribute, you can calculate the index page hit ratio for the buffer pool as follows:  $1 - (\text{buffer pool index physical reads} / \text{buffer pool index logical reads})$  If the hit ratio is low, increasing the number of buffer pool pages might improve performance.

- The type is int.
- The unit is requests.

#### Buffer Pool Avg Data Page Read per Async Req

- The average number of pages read for each asynchronous request. This value is derived by dividing the value of the Pool Async Data Reads attribute by the value of the Pool Async Data Read Reqs attribute.
- The type is int.
- The unit is pages.

#### Buffer Pool Pool Sync Index Reads

- The number of index pages read synchronously (that is, physical index page reads that were performed by database manager agents) into the buffer pool. This value is derived by subtracting the value of the Pool Async Index Reads attribute from Pool Index Physical Reads attribute. By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working.
- The type is int.
- The unit is pages.

#### Buffer Pool Pool Data to Estore

- Number of buffer pool data pages copied to extended storage. Pages are copied from the buffer pool to extended storage when they are selected as victim pages. As a result of the copying process, there is sufficient space for new pages in the buffer pool.
- The type is int.
- The unit is pages.

#### Buffer Pool Avg Sync Read Time

- The average elapsed time used to perform a synchronous read. This value is derived by dividing the value of the Pool Sync Read Time attribute by the value of the Pool Sync Read attribute. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device.
- The type is int.
- The unit is ?.

#### Buffer Pool Direct Write Time

- The elapsed time (in milliseconds) required to perform the direct writes. Use the following formula to calculate the average direct write time per sector:  $\text{direct write time} / \text{direct writes to database A}$  high average time might indicate an I/O conflict.
- The type is int.
- The unit is milliseconds.

#### Buffer Pool Pool Sync Read Time

- The elapsed time used to perform all synchronous reads. This value is derived by subtracting the value of the Pool Async Read Time attribute from the value of the Pool Read Time attribute. Use this attribute to understand the I/O work being performed.
- The type is int.
- The unit is ?.

#### Buffer Pool Pool Sync Write Time

- The total elapsed time used to perform all synchronous writes. This value is derived by subtracting the value of the Pool Async Write Time attribute from the value of the Pool Write Time attribute.
- The type is int.
- The unit is ?.

#### Buffer Pool Files Closed

- The total number of database files closed. The database manager opens files for reading and writing into and out of the buffer pool. The maximum number of database files open by an application at any time is controlled by the MAXFILOP configuration parameter. If the maximum is reached, one file is closed before the new file is opened. Note that the actual number of files opened might not equal the number of files closed. The following value is also valid:
- The type is int.
- The unit is files.

#### Buffer Pool Pool Data Writes

- The number of times a buffer pool data page was physically written to disk. A buffer pool data page is written to disk for the following reasons:
  - To free a page in the buffer pool so another page can be read
  - To flush the buffer pool.

If a buffer pool data page is written to disk for a high percentage of Buffer Pool Data Physical Reads, performance might improve by increasing the number of buffer pool pages available for the database.

- The type is int.
- The unit is writes.

#### Buffer Pool Pool Data from Estore

- Number of buffer pool data pages copied from extended storage. Required pages are copied from extended storage to the buffer pool. The copy process might incur the cost of connecting to the shared memory segment, but it saves the cost of a disk read. The following value is also valid:
- The type is int.
- The unit is pages.



#### Buffer Pool Pool Async Data Read Reqs

- The number of asynchronous read requests. To calculate the average number of data pages read per asynchronous request, use the following formula:  $\text{buffer pool asynchronous data reads} / \text{buffer pool asynchronous read requests}$  This average can help to determine the amount of asynchronous I/O in each interaction with the prefetcher.
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Async Read Time

- The total elapsed time spent reading by database manager prefetchers. Use this attribute to calculate the elapsed time for synchronous reading, using the following formula:  $\text{total buffer pool physical read time} - \text{buffer pool synchronous read time}$  You can also use this attribute to calculate the average asynchronous read time using the following formula:  $\text{buffer pool asynchronous read time} / \text{buffer pool asynchronous data reads}$  These calculations can be used to understand the I/O work being performed.
- The type is int.
- The unit is ?.

#### Buffer Pool Pool Async Data Reads

- The number of pages read asynchronously into the buffer pool. Use this attribute with the Buffer Pool Data Physical Reads attribute to calculate the number of physical reads that were performed synchronously (that is, physical data page reads that were performed by database manager agents). Use the following formula:  $\text{buffer pool data physical reads} - \text{buffer pool synchronous data reads}$  By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working.
- The type is int.
- The unit is pages.

#### Buffer Pool Pool Sync Index Writes

- The number of physical index write requests that were performed synchronously (that is, physical index page writes that were performed by database manager agents). This value is derived by subtracting the value of the Pool Async Index Writes attribute from the value of the Pool Index Writes attribute. By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing.
- The type is int.
- The unit is requests.

#### Buffer Pool Direct Read Time

- The elapsed time (in milliseconds) required to perform the direct reads. Use the following formula to calculate the average direct read time per sector:  $\text{direct read time} / \text{direct reads from database A}$  A high average time might indicate an I/O conflict.
- The type is int.
- The unit is milliseconds.

#### Buffer Pool Avg Pool Write Time

- The average elapsed time for a write request. This value is derived by dividing the value of the Pool Write Time attribute by the value of the Pool Total Writes attribute.
- The type is int.
- The unit is ?.

#### Buffer Pool Prefetch Ratio

- The percentage of asynchronous read operations that the prefetcher performed for sequential scans.
- The type is int.
- The unit is percent.

#### Buffer Pool Pool Index P Reads

- The number of physical read requests to get index pages into the buffer pool. see the Pool Index Logical Reads attribute for information about how to use this element.
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Sync Read

- The total number of synchronous reads. This value is derived by adding the values of the Pool Sync Data Reads and Pool Sync Index Reads attributes.
- The type is int.
- The unit is reads.

#### Buffer Pool Pool Index to Estore

- Number of buffer pool index pages copied to extended storage. Pages are copied from the buffer pool to extended storage when they are selected as victim pages. As a result of the copying process, there is sufficient space for new pages in the buffer pool.
- The type is int.
- The unit is pages.

#### Buffer Pool Pool Sync Data Reads

- The number of physical data page reads that were performed by database manager agents. This value is derived by subtracting the value of the Pool Async Data Reads attribute from the Pool Data Physical Reads attribute. By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working.
- The type is int.
- The unit is reads.

#### Buffer Pool Pool Async Index Reads

- The number of index pages read asynchronously into the buffer pool by a prefetcher. Asynchronous reads are performed by database manager prefetchers. Use this attribute with the Buffer Pool Index Physical Reads attribute to calculate the number of physical reads that were performed synchronously (that is, physical index page reads that were performed by database manager agents). Use the following formula: buffer pool index physical reads - buffer pool asynchronous index reads. By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working.
- The type is int.
- The unit is reads.

#### Buffer Pool Pool Total Writes

- The total number of write requests. This attribute is the total of the Pool Data Writes and Pool Index Writes attributes. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal.
- The type is int.
- The unit is writes.

#### Buffer Pool Avg Direct Read Time

- The average elapsed time for a direct read request. This value is calculated by dividing the value of the Direct Read Time attribute by the value of the the Direct Reads attribute. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device.
- The type is int.
- The unit is ?.

#### Buffer Pool Pool Total Reads

- The total number of read requests that required I/O to get data pages and index pages into the buffer pool. This attribute is the total of the Pool Data Physical Reads and Pool Index Physical Reads attributes. Values that are greater than or equal to 9223372036854775807 are indicated with the Value Exceeds Maximum text in the portal. The following value is valid:
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Async Data Writes

- The number of times a buffer pool data page was physically written to disk by an asynchronous page cleaner or by a prefetcher. A prefetcher might have written dirty pages to disk to make space for the pages being prefetched. Use this attribute with the Buffer Pool Data Writes attribute to calculate the number of physical write requests that were performed synchronously (that is, physical data page writes that were performed by database manager agents). Use the following formula:  $\text{buffer pool data writes} - \text{buffer pool asynchronous data writes}$  By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing.
- The type is int.
- The unit is writes.

#### Buffer Pool Pool Data P Reads

- The number of read requests that required I/O to get data pages into the buffer pool.
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Async Index Writes

- The number of times a buffer pool index page was physically written to disk by an asynchronous page cleaner or a prefetcher. A prefetcher might have written dirty pages to disk to make space for the pages being prefetched. Use this attribute with the Buffer Pool Index Writes attribute to calculate the number of physical index write requests that were performed synchronously. That is, physical index page writes that were performed by database manager agents. Use the following formula:  $\text{buffer pool index writes} - \text{buffer pool asynchronous index writes}$  By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing.
- The type is int.
- The unit is writes.

#### Buffer Pool Direct Writes

- The number of write operations that do not use the buffer pool. Use the following formula to calculate the average number of sectors that are written by a direct write:  $\text{direct writes to database} / \text{direct write requests}$  When using system monitors to track I/O, this data attribute helps to distinguish database I/O from non-database I/O on the device.
- The type is int.

- The unit is writes.

#### Buffer Pool Pool Write Time

- The total amount of time spent physically writing data or index pages from the buffer pool to disk. Use this attribute with the Buffer Pool Data Writes and Buffer Pool Index Writes attributes to calculate the average page-write time. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device. The following value is valid:
- The type is int.
- The unit is ?.

#### Buffer Pool Application Direct Read Reqs

- The number of requests to perform a direct read of one or more sectors of data. Use the following formula to calculate the average number of sectors that are read by a direct read:  $\text{direct reads from database} / \text{direct read requests}$  The following value is also valid:
- The type is int.
- The unit is requests.

#### Buffer Pool Pool Read Time

- The total amount of elapsed time spent processing read requests that caused data or index pages to be physically read from buffer pool to disk. Use this attribute with the Buffer Pool Data Physical Reads and Buffer Pool Index Physical Reads attributes to calculate the average page-read time. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device. The following value is also valid:
- The type is int.
- The unit is ?.

### **Component: Customized SQL Definitions**

information for customized SQL statements, such as the name of the definition file, the last modified time of the definition file, SQL ID and SQL content.

#### **Dimensions**

##### Custom SQL Definition Customized Definition File

- The location of the definition file for customized SQL, which is a key attribute.
- The type is string. This is a key dimension.

##### Custom SQL Definition Node Name

- The managed system name of the agent. For new installations of version 7. 1, the format is instanceid:hostname:UD for all operating systems.
- The type is string.

##### Custom SQL Definition SQL Content

- The SQL content that is defined in the definition file. The carriage return is replaced by a blank. The shown text is limited to 512 bytes.
- The type is string.

##### Custom SQL Definition SQL ID

- The SQL ID that is defined in the definition file.
- The type is string. This is a key dimension.

#### Custom SQL Definition Last Modified Time

- The last time that the definition file was modified.
- The type is timestamp.

#### Metrics

##### Custom SQL Definition Time Stamp

- The local time at the agent when the data was collected.
- The type is timestamp.
- The unit is unspecified.

#### Component: System Resources

Information about the OS environment in which the DB2 instance is running. This data set is only available for DB2 Version 9. 5 and later.

#### Dimensions

##### System Resources OS Version

- The version number of the operating system.
- The type is string.

##### System Resources Host Name

- The name of the host that owns the resources.
- The type is string.

##### System Resources OS Level

- The maintenance level of the current version and release. For example, LINUX: 2. 4. 9, level = 9.
- The type is string.

##### System Resources Operating System Name

- The full name of the operating system.
- The type is string.

##### System Resources Machine Identification

- The machine hardware identification.
- The type is string.

##### System Resources OS Release

- The release of the operating system. For example, AIX: 4. 3 release = 3.
- The type is string.

##### System Resources Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string.

#### Metrics

##### System Resources Total Swap Memory (Superseded)

- The total amount of the swap memory on the system.
- The type is int.

- The unit is megabytes.

#### System Resources Free Swap Memory

- The total amount of the free swap memory on the system.
- The type is int.
- The unit is megabytes.

#### System Resources Total Virtual Memory

- The total amount of the virtual memory on the system.
- The type is int.
- The unit is megabytes.

#### System Resources Pct of Physical Memory Used

- The percentage of the physical memory that is used on the system.
- The type is double.
- The unit is percent.

#### System Resources Pct of Virtual Memory Used

- The percentage of the virtual memory that is used on the system.
- The type is double.
- The unit is percent.

#### System Resources Pct of CPU Used

- The percentage of the CPU that is used on the system.
- The type is int.
- The unit is percent.

#### System Resources Total Swap Memory

- The total amount of swap memory on the system.
- The type is int.
- The unit is megabytes.

#### System Resources Pct of Swap Memory Used (Superseded)

- The percentage of the swap memory that is used on the system.
- The type is int.
- The unit is percent.

#### System Resources Free Physical Memory

- The amount of free physical memory on the system.
- The type is int.
- The unit is megabytes.

#### System Resources Total Virtual Memory (Superseded)

- The total amount of the virtual memory on the system.
- The type is int.
- The unit is megabytes.

#### System Resources Total Physical Memory

- The total amount of the physical memory on the system.
- The type is int.
- The unit is megabytes.

#### System Resources Pct of Virtual Memory Used (Superseded)

- The percentage of the virtual memory that is used on the system.
- The type is int.
- The unit is percent.

#### System Resources Pct of Physical Memory Used (Superseded)

- The percentage of the physical memory that is used on the system.
- The type is int.
- The unit is percent.

#### System Resources Free Virtual Memory

- The total amount of the free virtual memory on the system.
- The type is int.
- The unit is megabytes.

#### System Resources Pct of Swap Memory Used

- The percentage of the swap memory that is used on the system.
- The type is double.
- The unit is percent.

### **Component: Slow SQL Statements**

Information about slow SQL Statements.

#### **Dimensions**

##### Slow SQL DB Partition

- The DB2 database partition node number. DB2 partition numbers range from 0 to 999. The 'Aggregated' and 'Current' values can be used within a query or situation filter. If no db partition filter is specified, then a row of data will be returned for each database partition. If a db partition filter is used with the 'Aggregated' value, then only aggregated partition data will be returned. Historical data collection will include both aggregated and individual partition attribute data.
- The type is int. This is a key dimension.

##### Slow SQL Node Name

- The origin node of db2 agent.
- The type is string.

##### Slow SQL Statement Type

- The type of the SQL statement, such as static or dynamic.
- The type is string.

##### Slow SQL Statement Text

- The query for the SQL statement.
- The type is string. This is a key dimension.

#### Slow SQL Active State

- The state of the SQL statement.
- The type is string.

#### Slow SQL Duration

- The duration of executing the SQL statement.
- The type is string.

#### Slow SQL Stmt Start Timestamp

- The start time of the SQL statement.
- The type is timestamp.

#### Slow SQL DB Name

- The name of the monitored database.
- The type is string. This is a key dimension.

#### Slow SQL Executable ID

- The unique identifier for the SQL statement.
- The type is string. This is a key dimension.

### Metrics

#### Slow SQL Lock wait

- The total number of times that applications or connections waited for locks while executing the SQL Statement.
- The type is int.
- The unit is waits.

### Component: Custom SQL Execution

Information about the results of customized SQL statement executions, including five string columns, five number columns, and two date and time columns.

#### Dimensions

##### Custom SQL Fourth String Column Name

- The name of the fourth string type column in the result of the customized SQL statement execution.
- The type is string.

##### Custom SQL Second Date Value

- The second date time value in the result of the customized SQL statement execution.
- The type is timestamp.

##### Custom SQL Fifth String Column Name

- The name of the fifth string type column in the result of the customized SQL statement execution.
- The type is string.

##### Custom SQL Second String Column Name



- The name of the second string type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Third String Column Name

- The name of the third string type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL First String Column Name

- The name of the first string type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Fifth Number Column Name

- The name of the fifth number type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Fourth Number Column Name

- The name of the fourth number type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Third Number Column Name

- The name of the third number type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Second Number Column Name

- The name of the second number type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Status SQL ID

- The SQL ID that is defined in the definition file.
- The type is string. This is a key dimension.

#### Custom SQL First Number Column Name

- The name of the first number type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Status Node Name

- The managed system name of the agent. For new installations of version 7. 1, the format is instanceid:hostname:UD for all operating systems.
- The type is string.

#### Custom SQL Status Last Execution Error Message

- The error message returned by DB2 for the last SQL execution, which has a maximum length of 256 characters.
- The type is string.

#### Custom SQL Status Last Execution Time

- The last date and time when the SQL is executed.
- The type is timestamp.

#### Custom SQL Fifth String Value

- The fifth string value in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Fourth String Value

- The fourth string value in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Third String Value

- The third string value in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Second String Value

- The second string value in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL First String Value

- The first string value in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Status DB Alias

- The alias name of the database on which the SQL Statement associated with the SQL ID is executed, which is a key attribute.
- The type is string. This is a key dimension.

#### Custom SQL First Date Column Name

- The name of the first date time type column in the result of the customized SQL statement execution.
- The type is string.

#### Custom SQL Status SQL State

- The SQL STATE returned by DB2 for the last SQL execution, which has a length of 10 characters.
- The type is string.

#### Custom SQL DB Alias Filter Name

- The Database alias filter name that can be defined as:
  - The character ( \* ), is required if you want to execute the SQL statement associated with the SQL ID on all the databases of the DB2 server excluding all HADR standby databases.
  - A database alias, this is required if you want to execute the SQL statement associated with the SQL ID on a specific database.

If the database filter alias name contains blank spaces at the beginning and at the end, the blank spaces at the end are trimmed.

- The type is string.

Custom SQL Second Date Column Name

- The name of the second date time type column in the result of the customized SQL statement execution.
- The type is string.

Custom SQL First Date Value

- The first date time value in the result of the customized SQL statement execution.
- The type is timestamp.

## **Metrics**

Custom SQL First Number Value

- The first number value in the result of the customized SQL statement execution.
- The type is int.
- The unit is unspecified.

Custom SQL Second Number Value

- The second number value in the result of the customized SQL statement execution.
- The type is int.
- The unit is unspecified.

Custom SQL Third Number Value

- The third number value in the result of the customized SQL statement execution.
- The type is int.
- The unit is unspecified.

Custom SQL Fourth Number Value

- The fourth number value in the result of the customized SQL statement execution.
- The type is int.
- The unit is unspecified.

Custom SQL Fifth Number Value

- The fifth number value in the result of the customized SQL statement execution.
- The type is int.
- The unit is unspecified.

Custom SQL Status Last Execution Error Code

- The native error code returned by DB2 for the last SQL execution.
- The type is int.
- The unit is unspecified.

## **Component: Application Activities**

Information about application activities.

## Dimensions

### Locking Conflict Tablespace Name (Unicode)

- The name of the tablespace against which the application currently holds a lock (Unicode). The value format is a text string with a maximum of 60 bytes.
- The type is string.

### Application Group Stmt Start

- The string date and time that the most recent SQL statement operation started. The value format is CYYMMDDHHMMSSmmm. Use this attribute with the Statement Stop attribute to calculate the elapsed execution time for the statement operation.
- The type is string.

### Locking Conflict Client DB Alias

- The alias defined within the database manager where the database connection request originated. The value format is a text string with a maximum of 20 characters. Use to identify the actual database that the application is accessing. The mapping between this name and Database Name can be done by using the database directories at the client node and the database manager server node. Because different database aliases can have different authentication types, this attribute can also help you determine the authentication type.
- The type is string.

### Application Lock Wait Start Time

- The date and time that the application started waiting to obtain a lock on the object that is currently locked by another application.
- The type is timestamp.

### Application Group Agent Sys CPU Time

- The total system CPU time (in seconds) that the database manager agent process spent executing database manager code. This element includes CPU time for both SQL and non-SQL statements, and CPU time for any unfenced user-defined functions (UDFs).
- The type is string.

### Application Auth ID

- The authorization ID of the user who invoked the application that is being monitored. On a DB2 Connect gateway node, this ID is the user authorization ID on the host. The value format is a text string with a maximum of 60 bytes. Use this attribute to determine who invoked the application.
- The type is string.

### Locking Conflict Lock Escalation

- An indicator of whether a lock request was made as part of a lock escalation.
  - No
  - Yes

Use this attribute to better understand the cause of deadlocks. If deadlocks occur that involve applications doing lock escalation, you might want to increase the amount of lock memory or change the percentage of locks that any one application can request.

- The type is string.

### Application Group Stmt Stop

- The string date and time that the most recent SQL statement operation stopped. If the statement is still running, this field is 0 (zero). Use this attribute with the Statement Start attribute to calculate the elapsed execution time for the statement operation.
- The type is string.

#### Locking Conflict Client DB Alias (Unicode)

- The alias defined within the database manager where the database connection request originated (Unicode). The value format is a text string with a maximum of 60 bytes.
- The type is string.

#### Application Stmt Text

- The text of the dynamic SQL statement. For application snapshots, the statement text helps you identify what the application was executing when the snapshot was taken, or most recently processed if no statement was being processed at the time the snapshot was taken. For dynamic SQL statements, this attribute identifies the SQL text associated with a package. The value format is a text string with a maximum of 2000 bytes.
- The type is string.

#### Application Agent ID Holding Lock

- The application handle of the agent holding a lock for which this application is waiting. The value format is an integer. The lock monitor group must be turned on to obtain this information.
- The type is int.

#### Locking Conflict Appl ID (Unicode)

- The identifier generated when the application connects to the database at the database manager or when DDCS receives a request to connect to a DRDA database (Unicode).
- The type is string.

#### Application Table Schema

- The schema of the table the application is waiting to lock. The value format is a text string with a maximum of 60 bytes. Along with the Table Name attribute, this attribute can help to determine the source of contention for resources.
- The type is string.

#### Application Creator

- The authorization ID of the user that precompiled the application (Unicode). The value format is a text string with a maximum of 60 bytes. Use this attribute to help identify the SQL statement that is processing, with the CREATOR column of the package section information in the catalogs.
- The type is string.

#### Application Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string.

#### Application Client Protocol

- The communication protocol that the client application is using to communicate with the server. The value format is a text string with a maximum of 12 characters. Use this attribute for troubleshooting of remote applications.
- The type is string.

#### Application Group UOW Stop Time

- The string date and time that the most recent unit of work completed, which occurs when database changes are committed or rolled back. The value format is CYYMMDDHHMMss.
- The type is string.

#### Locking Conflict Lock Wait Start Timestamp

- The date and time that the application started waiting to obtain a lock on the object that is currently locked by another application.
- The type is timestamp.

#### Application Appl ID Holding Lock

- The application ID of the application that is holding a lock on the object that this application is waiting to obtain (Unicode). The value format is a text string with a maximum of 96 bytes.
- The type is string.

#### Application Conn Complete Timestamp

- The date and time that a connection request was granted.
- The type is timestamp.

#### Application Client Prdid

- The product and version identifier for the software on the client. The value format is a text string with a maximum of 20 characters. For example: SQL06010 .
- The type is string.

#### Application Appl ID

- The identifier generated when the application connects to the database at the database manager or when DDCS receives a request to connect to a DRDA database. The value format is a text string, with a maximum of 32 characters.
- The type is string. This is a key dimension.

#### Application Client Platform

- The operating system on which the client application is running. Use this attribute to analyze problems for remote applications. The value format is a text string with a maximum of 20 characters.
- The type is string.

#### Application Lock Mode

- The type of lock being held. Use the lock mode to determine the source of contention for resources. The value format is a text string with a maximum of 32 characters.
- The type is string.

#### Application Group UOW Comp Status

- The completion status of the previous UOW (unit of work). Use this attribute to determine if the unit of work ended due to a deadlock or an abnormal termination.
- The type is string.

#### Locking Conflict Snapshot Time

- The string date and time when the database system monitor information was collected. Use this attribute to help relate data chronologically if you are saving the results in a file or database for ongoing analysis. The timestamp value is formatted as a date and time string. The internal timestamp value that is stored in the database is in the format cYYMMDDhhmmss000.

- The type is string.

#### Application Lock Object Type

- The type of object against which the application holds a lock (for object-lock-level information), or the type of object for which the application is waiting to obtain a lock (for application-level and deadlock-level information). The value format is a text string with a maximum of 16 characters.
- The type is string.

#### Application Corr Token

- The DRDA AS correlation token. The value format is a text string with a maximum of 96 bytes.
- The type is string.

#### Application Execution ID

- The ID that the user specified when logging in to the operating system. This ID is distinct from the Authorization ID, which the user specifies when connecting to the database. The value format is a text string with a maximum of 60 bytes. Use this attribute to determine the operating system user ID of the individual running the monitored application.
- The type is string.

#### Locking Conflict Auth ID (Unicode)

- The authorization ID of the user who invoked the application that is being monitored (Unicode). On a DB2 Connect gateway node, this is the user authorization ID on the host. The value format is a text string with a maximum of 20 bytes.
- The type is string.

#### Application Group Database Manager Agent User CPU Time

- The total CPU time (in microseconds) that the database manager agent process used. This counter includes time spent on both SQL and non-SQL statements, and any fenced user-defined functions (UDFs) or stored procedures issued by the application. System CPU represents the time spent in system calls. User CPU represents time spent executing database manager code. The value format is a text string with a maximum of 10 characters. Use this attribute with the other CPU-time related attributes to help you identify applications or queries that consume large amounts of CPU time.
- The type is string.

#### Application Stmt Type

- The type of SQL statement processed. The value format is a text string with a maximum of 32 characters.
- The type is string.

#### Application Client PID

- The process ID of the client application that made the connection to the database. The value format is an integer. Use this attribute to correlate monitor information such as CPU and I/O time to your client application. If a DRDA AS connection is used, this element is set to 0.
- The type is int.

#### Application DB Name

- The real name of the database for which information is collected or to which the application is connected. This name was given to the database when it was created. The value format is a

simple text string with a maximum of 60 bytes. Use this attribute to identify the specific database to which the data applies.

- The type is string.

#### Application Group Prev UOW Stop Time

- The string date and time that the unit of work completed. The value format is CYYMMDDHHMMSSmmm. Use this attribute with the UOW Stop Time attribute to calculate the total elapsed time between COMMIT/ROLLBACK points, and with the UOW Start Time attribute to calculate the time spent in the application between units of work.
- The type is string.

#### Locking Conflict Table Schema (Unicode)

- The schema of the table against which the application is holding a lock (Unicode). The value format is a text string with a maximum of 60 bytes.
- The type is string.

#### Locking Conflict Table Name (Unicode)

- The name of the table against which the application is holding locks (Unicode). The value format is a text string with a maximum of 60 bytes.
- The type is string.

#### Application Appl Status

- The status of the application being monitored. This attribute can help you diagnose potential application problems. The value format is a text string with a maximum of 64 characters.
- The type is string.

#### Locking Conflict Status Change Time

- The string date and time the application entered its current status. The value format is CYYMMDDHHMMSSmmm. Use this attribute to determine how long an application has been in its current status. If the application status remains unchanged for a long period of time, the application might have a problem.
- The type is string.

#### Locking Conflict Appl Name (Unicode)

- The name of the application running at the client as it is known to the database manager or DB2 Connect (Unicode). The value format is a text string, with a maximum of 60 bytes.
- The type is string. This is a key dimension.

#### Application Package Name

- The name of the package that contains the SQL statement currently executing. The value format is a text string with a maximum of 60 bytes.
- The type is string.

#### Application Table Name

- The name of the table the application is waiting to lock. The value format is a text string with a maximum of 60 bytes. Use this attribute with the Table Schema attribute to determine the source of contention for resources.
- The type is string.

#### Application Group Section Number



- The internal section number in the package for the SQL statement currently processing or most recently processed.
- The type is int.

#### Application Group UOW Start Time

- The string date and time that the unit of work first required database resources. This resource requirement occurs at the first SQL statement execution for the unit of work. The value format is CYYMMDDHHMMsss. Use this attribute with the UOW Stop Time attribute to calculate the total elapsed time of the unit of work and with the Previous Unit of Work Completion Timestamp attribute to calculate the time spent in the application between units of work.
- The type is string.

#### Application Agent ID

- The application handle, which is a system-wide unique ID for the application. The value format is an integer. On multi-node systems, where a database is partitioned, this ID is the same on every node where the application might make a secondary connection.
- The type is int.

#### Application Tablespace Name

- The name of the tablespace that the application is waiting to lock. The value format is a text string with a maximum of 60 bytes. This attribute can help you to determine the source of contention for resources.
- The type is string.

#### Locking Conflict Appl ID Holding Lock (Unicode)

- The application ID of the application that is holding a lock on the object that this application is waiting to obtain (Unicode). The value format is a text string with a maximum of 96 bytes.
- The type is string.

#### Application Country Code

- The country code of the client application. The value format is an integer.
- The type is int.

#### Application Appl Name

- The name of the application that is connected to the database. The value format is a text string, with a maximum of 60 bytes. For example: \*Local. db2inst1. 990212202018 .
- The type is string. This is a key dimension.

#### Application DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If a db partition filter is not specified, data is returned for the current database partition. If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range.
- The type is string. This is a key dimension.

#### Application Appl Conn Timestamp

- The date and time that an application started a connection request.
- The type is timestamp.

#### Locking Conflict Status Change Timestamp

- The date and time the application entered its current status.
- The type is timestamp.

#### Application Cursor Name

- The name of the cursor corresponding to this SQL statement. The value format is a text string with a maximum of 60 bytes.
- The type is string.

#### Application Stmt Operation

- The statement operation currently being processed or most recently processed (if none is currently running). The value format is a text string with a maximum of 20 characters.
- The type is string.

#### Locking Conflict Codepage ID

- The codepage or CCSID at the node where the application started. For snapshot monitor data, this is the code page at the node where the monitored application started. Use this attribute to analyze problems for remote applications. By using this information, you can ensure that data conversion is supported between the application code page and the database code page (or for DRDA host databases, the host CCSID).
- The type is int.

### Metrics

#### Application Pool Data to Estore

- Number of buffer pool data pages copied to extended storage. The value format is an integer.
- The type is int.
- The unit is pages.

#### Application Rows Read

- The number of rows read from the table. The value format is an integer. This attribute helps to identify tables with heavy usage for which you might want to create additional indexes.
- The type is int.
- The unit is reads.

#### Application Direct Writes

- The number of write operations that do not use the buffer pool. The value format is an integer.
- The type is int.
- The unit is writes.

#### Application Group Lock Waits

- The total number of times that applications or connections waited for locks. At the database level, the lock waits value is the total number of times that applications waited for locks within this database. At the application-connection level, the lock waits value is the total number of times that this connection requested a lock but waited because another connection was already holding a lock on the data. Use this attribute with the Lock Wait Time attribute to calculate, at the database level, the average wait time for a lock. This calculation can be performed at either the database or the application-connection level. If the average lock wait time is high, look for applications that hold many locks, or have lock escalations, with a focus on tuning your applications to improve concurrency, if appropriate. If escalations are the reason for a high

average lock wait time, the values of one or both of the LOCKLIST and MAXLOCKS configuration parameters might be too low.

- The type is int.
- The unit is waits.

#### Application Group Locks Held

- The number of locks currently held. If the monitor information is at the database level, this number represents the total number of locks currently held by all applications in the database. If the information is at the application level, this number represents the total number of locks currently held by all agents for the application.
- The type is int.
- The unit is locks.

#### Application Pool Hit Ratio

- The buffer pool hit ratio (as a percentage). The value format is an integer. The sum of the Pool Data Logical Reads and Pool Index Logical Reads attributes is divided by the value of the Pool Total Reads attribute to derive the percentage.
- The type is double.
- The unit is percent.

#### Application Pool Data P Reads

- The number of read requests that required I/O to get data pages into the buffer pool. The value format is an integer.
- The type is int.
- The unit is requests.

#### Application Avg Pool Write Time

- The average elapsed time for a write request. The value format is an integer.
- The type is int.
- The unit is ?.

#### Application UOW Lock Wait Time

- The time the UOW (unit of work) waited on locks (in seconds).
- The type is int.
- The unit is seconds.

#### Application Direct Read Time

- The elapsed time (in milliseconds) required to perform the direct reads. The value format is an integer.
- The type is int.
- The unit is milliseconds.

#### Application Rows Selected

- The number of rows that have been selected and returned to the application. The value format is an integer. Use this attribute to gain insight into the current level of activity within the database manager.
- The type is int.
- The unit is selects.

#### Application Cat Cache Overflows

- The number of times that an insert into the catalog cache failed because the catalog cache was full. The value format is an integer. If the catalog cache overflows value is large, the catalog cache might be too small for the workload. Increasing the size of the catalog cache might improve its performance. If the workload includes transactions that compile a large number of SQL statements referencing many tables, views, and aliases in a single unit of work, compiling fewer SQL statements in a single transaction might improve the performance of the catalog cache. Or if the workload includes the binding of packages containing many SQL statements referencing many tables, views or aliases, you might want to split the packages so that they include fewer SQL statements to improve performance.
- The type is int.
- The unit is overflows.

#### Application Cat Cache Inserts

- The number of times that the system tried to insert table descriptor information into the catalog cache. The value format is an integer. Table descriptor information is inserted into the cache following a failed lookup to the catalog cache while processing a table, view, or alias reference in an SQL statement. The catalog cache inserts value includes attempts to insert table descriptor information that fail due to catalog cache overflow and heap full conditions.
- The type is int.
- The unit is inserts.

#### Application Rows Deleted

- The number of row deletions attempted. The value format is an integer. Use this attribute to gain insight into the current level of activity within the database manager.
- The type is int.
- The unit is deletes.

#### Application Pkg Cache Lookups

- The number of times that an application looked for a section or package in the package cache. The value format is an integer. At a database level, it indicates the overall number of references since the database was started, or monitor data was reset. Note that this counter includes the cases where the section is already loaded in the cache and when the section has to be loaded into the cache.
- The type is int.
- The unit is lookups.

#### Application Int Rollbacks

- The total number of rollbacks initiated internally by the database manager. The value format is an integer.
- The type is int.
- The unit is rollbacks.

#### Application Binds Precompiles

- The number of binds and precompiles attempted. The value format is an integer. Use this attribute to gain insight into the current level of activity within the database manager.
- The type is int.
- The unit is operations.

#### Application Lock Escals

- The number of times that locks have been escalated from several row locks to a table lock. A lock is escalated when the total number of locks held by an application reaches the maximum amount of lock list space available to the application, or the lock list space consumed by all applications is approaching the total lock list space. This data item includes a count of all lock escalations, including exclusive lock escalations. When an application reaches the maximum number of locks allowed and there are no more locks to escalate, the application uses space in the lock list that is allocated for other applications. When the entire lock list is full, an error occurs. The value format is an integer.
- The type is int.
- The unit is occurrences.

#### Application application Static SQL Stmts

- The number of static SQL statements that were attempted. The value format is an integer.
- The type is int.
- The unit is statements.

#### Application Open Local Curs Blk

- The number of local blocking cursors currently open for this application. The value format is an integer.
- The type is int.
- The unit is cursors.

#### Application Commit SQL Stmts

- The total number of SQL COMMIT statements that have been attempted. The value format is an integer. A small rate of change in this counter during the monitor period might indicate that applications are not doing frequent commits. The lack of frequent commits can lead to problems with logging and data concurrency. The following value is also valid:
- The type is int.
- The unit is commits.

#### Application Pool Data from Estore

- Number of buffer pool data pages copied from extended storage. The value format is an integer.
- The type is int.
- The unit is pages.

#### Application Avg Pool Read Time

- The average elapsed time for a read request. The value format is an integer.
- The type is int.
- The unit is ?.

#### Application Pkg Cache Hit Ratio

- The percentage of package sections that were found in the cache. The value format is an integer.
- The type is double.
- The unit is percent.

#### Application Open Local Curs

- The number of local cursors currently open for this application, including those cursors counted by Open Local Cursors with Blocking attribute. The value format is an integer.

- The type is int.
- The unit is cursors.

#### Application Hash Join Small Overflows

- The number of times that hash join data exceeded the available sort heap space by less than 10%. The value format is an integer. If this value and the value of the Hash Join Overflows attribute are high, you must consider increasing the sort heap threshold. If this value is greater than 10% of Hash Join Overflows, you must consider increasing the sort heap size.
- The type is int.
- The unit is occurrences.

#### Application X Lock Escals

- The number of times that locks have been escalated from several row locks to one exclusive table lock, or the number of times an exclusive lock on a row caused the table lock to become an exclusive lock. The value format is an integer. A lock is escalated when the total number of locks held by an application reaches the maximum amount of lock list space available to the application. The amount of lock list space available is determined by the LOCKLIST and MAXLOCKS configuration parameters. Other applications cannot access data held by an exclusive lock.
- The type is int.
- The unit is escalations.

#### Application Hash Join Overflows

- The number of times that hash join data exceeded the available sort heap space. The value format is an integer.
- The type is int.
- The unit is occurrences.

#### Application Avg Sort Time

- The average time that was elapsed to complete a sort operation. At the database or application level, the value for this attribute can indicate the performance issues with sorting. This attribute value is affected by the system load. The following value is also valid:
- The type is int.
- The unit is value.

#### Locking Conflict Locks Held

- The number of locks currently held. If the monitor information is at the database level, this is the total number of locks currently held by all applications in the database. If it is at the application level, this is the total number of locks currently held by all agents for the application. Usage of this attribute depends on the level of information being returned from the database system monitor. The following value is also valid:
- The type is int.
- The unit is locks.

#### Application Select SQL Stmts

- The number of SQL SELECT statements that were issued. The value format is an integer.
- The type is int.
- The unit is selects.

#### Application Pool Index Writes

- The number of times a buffer pool index page was physically written to disk. The value format is an integer. If a buffer pool index page is written to disk for a high percentage of the Pool Index Physical Reads, performance might improve by increasing the number of buffer pool pages available for the database. If all applications are updating the database, increasing the size of the buffer pool might have minimal impact on performance; most pages contain updated data that must be written to disk.
- The type is int.
- The unit is writes.

#### Application Int Rows Inserted

- The number of rows inserted into the database as a result of internal activity caused by triggers. The value format is an integer. This attribute can help to gain insight into the internal activity within the database manager. If this activity is high, you must evaluate the design to determine if you can alter it to reduce this activity.
- The type is int.
- The unit is rows.

#### Application Int Auto Rebinds

- The number of automatic rebinds (or recompiles) that have been attempted. The value format is an integer. Automatic rebinds are the internal binds the system performs when a package has been invalidated. Use this attribute to determine the level of database activity at the application or database level.
- The type is int.
- The unit is rebinds.

#### Application Rows Written

- The number of rows changed (inserted, deleted, or updated) in the table. The value format is an integer. A high value for table-level information indicates heavy usage of the table. If so, you might want to use the Run Statistics (RUNSTATS) utility to maintain efficiency of the packages used for this table.
- The type is int.
- The unit is changes.

#### Application Cat Cache Hit Ratio

- The percentage of catalog sections that are found in the cache. The value format is an integer.
- The type is double.
- The unit is percent.

#### Application Pool Index L Reads

- The number of logical read requests for index pages that have gone through the buffer pool. The value format is an integer.
- The type is int.
- The unit is requests.

#### Application Cat Cache Heap Full

- The number of times that an insert into the catalog cache failed because of a heap full condition in the database heap. The value format is an integer. The catalog cache draws its storage dynamically from the database heap. Even if the cache storage has not reached its limit, inserts into the catalog cache might fail due to a lack of space in the database heap. If the catalog cache heap full count is not zero, you can correct the insert failure condition by increasing the database heap size or by reducing the catalog cache size.

- The type is int.
- The unit is failures.

#### Application Sort Overflows

- The total number of sorts that ran out of sort heap space and might have required disk space for temporary storage. The value format is an integer. at the database or application level, use this element with the Total Sorts attribute. This attribute can help to determine the source of contention for resources.
- The type is int.
- The unit is sorts.

#### Application Cat Cache Lookups

- The number of times that the catalog cache was referenced to obtain table descriptor information. The value format is an integer.
- The type is int.
- The unit is lookups.

#### Application Locks Held

- The number of locks that are currently held. The value format is an integer.
- The type is int.
- The unit is locks.

#### Application Pool Read Time

- The total amount of elapsed time spent processing read requests that caused data or index pages to be physically read from disk to buffer pool. The value format is an integer.
- The type is int.
- The unit is ?.

#### Application Failed SQL Stmts

- The number of SQL statements that were attempted, but failed. The value format is an integer.
- The type is int.
- The unit is statements.

#### Application Int Deadlock Rollbacks

- The total number of forced rollbacks initiated by the database manager due to a deadlock. The value format is an integer. The database manager initiates a rollback for the current unit of work in an application that is experiencing a deadlock. This attribute shows the number of deadlocks that have been broken. It can indicate the possibility of concurrency problems. It is also important because internal rollbacks due to deadlocks can cause performance degradation.
- The type is int.
- The unit is rollbacks.

#### Application Direct Write Time

- The elapsed time (in milliseconds) required to perform the direct writes. The value format is an integer.
- The type is int.
- The unit is milliseconds.

#### Application Total Sorts



- The total number of sorts that have been issued. The value format is an integer. at the database or application level, use this value with the Sort Overflows attribute to calculate the percentage of sorts that need more heap space. You can also use it with the Total Sort Time attribute to calculate the average sort time. If the number of sort overflows is small with respect to the total sorts, increasing the sort heap size might have little impact on performance, unless this buffer size is increased substantially.
- The type is int.
- The unit is sorts.

#### Application Deadlocks

- The total number of deadlocks that have occurred. The value format is an integer. This attribute can indicate that applications are experiencing contention problems. To resolve the problem, determine in which applications (or application processes) the deadlocks are occurring. You can then modify the application to enable it to run concurrently. Some applications, however, might not be capable of running concurrently.
- The type is int.
- The unit is deadlocks.

#### Application Pool Total Writes

- The total number of write requests. The value format is an integer. This attribute is the total of the Pool Data Writes and Pool Index Writes attributes. Use this attribute to determine how busy the DB2 server is in terms of write I/O activity. Values that are greater than or equal to 2147483647 are indicated in the portal with the Value Exceeds Maximum text, and values that are smaller than -2147483648 are indicated with the Value Exceeds Minimum text.
- The type is int.
- The unit is requests.

#### Application Rej Curs Blk

- The number of times that a request for an I/O block at the server was rejected and the request was converted to non-blocked I/O. If there are many cursors blocking data, the communication heap might become full. The value format is an integer. When this heap is full, I/O blocks are not allocated for blocking cursors; however, an error condition does not alert you to this condition. If cursors are unable to block data, performance can be affected adversely.
- The type is int.
- The unit is occurrences.

#### Application Int Rows Deleted

- The number of rows deleted from the database as a result of internal activity. The value format is an integer. This attribute can help to gain insight into internal activity within the database manager. If this activity is high, you must evaluate the table design to determine if the referential constraints or triggers that you defined on the database are necessary.
- The type is int.
- The unit is rows.

#### Locking Conflict Lock Wait Time

- The total elapsed time (in milliseconds) that a lock was waited for. At the database level, this is the total amount of elapsed time that all applications were waiting for a lock within this database. At the application-connection and transaction levels, this is the total amount of elapsed time that this connection or transaction has waited for a lock to be granted. This attribute might be used with the Lock Waits attribute to calculate the average wait time for a lock. This calculation can be performed at either the database or the application-connection level. The following value is also valid:

- The type is int.
- The unit is milliseconds.

#### Application Pool Index to Estore

- Number of buffer pool index pages copied to extended storage. The value format is an integer. Pages are copied from the buffer pool to extended storage when they are selected as victim pages. As a result of the copying process, there is sufficient space for new pages in the buffer pool.
- The type is int.
- The unit is pages.

#### Application Pkg Cache Inserts

- The total number of times that a requested section was not available for use and had to be loaded into the package cache. The value format is an integer. This count includes any implicit prepares performed by the system.
- The type is int.
- The unit is inserts.

#### Application Pool Total Reads

- The total number of read requests that required I/O to get data pages and index pages into the buffer pool. The value format is an integer. This attribute is the total of the Pool Data Physical Reads and Pool Index Physical Reads attributes. Use this attribute to determine how busy the DB2 server is in terms of I/O activity. Values that are greater than or equal to 2147483647 are indicated in the portal with the Value Exceeds Maximum text, and values that are smaller than -2147483648 are indicated with the Value Exceeds Minimum text.
- The type is int.
- The unit is requests.

#### Application Appl Idle Time

- The number of seconds since an application issued a request to the server. The value format is an integer.
- The type is int.
- The unit is seconds.

#### Application Direct Read Reqs

- The number of requests to perform a direct read of one or more sectors of data. The value format is an integer.
- The type is int.
- The unit is requests.

#### Application Avg Lock Wait Time

- The average elapsed time (in milliseconds) that was spent waiting for a lock. The value format is an integer. If the average lock wait time is high, you must look for applications that hold many locks, or have lock escalations, with a focus on tuning your applications to improve concurrency, if appropriate. If escalations are the reason for a high average lock wait time, the values of one or both of the LOCKLIST and MAXLOCKS configuration parameters might be too low.
- The type is int.
- The unit is milliseconds.

#### Application Lock Timeouts

- The number of times that a request to lock an object time out instead of being granted. The value format is an integer.
- The type is int.
- The unit is timeouts.

#### Application Group Prefetch Wait Time

- The time an application spent waiting for an I/O server or prefetcher to finish loading pages into the buffer pool. This attribute can be used to experiment with changing the number of I/O servers and the I/O server sizes.
- The type is int.
- The unit is ?.

#### Application Pool Index P Reads

- The number of physical read requests to get index pages into the buffer pool. The value format is an integer.
- The type is int.
- The unit is requests.

#### Application Pool Data L Reads

- The number of logical read requests for data pages that have gone through the buffer pool. The value format is an integer. This count includes accesses to data that is already in the buffer pool when the database manager needs to process the page or read into the buffer pool before the database manager can process the page.
- The type is int.
- The unit is requests.

#### Application Query Cost Estimate

- Estimated cost, in timerons, for a query, as determined by the SQL compiler. The value format is an integer. This attribute allows correlation of actual runtime values with the compile-time estimates.
- The type is int.
- The unit is timerons.

#### Application Int Rows Updated

- The number of rows updated from the database as a result of internal activity. The value format is an integer. This attribute can help to gain insight into internal activity within the database manager. If this activity is high, you must evaluate the table design to determine if the referential constraints that you defined on the database are necessary.
- The type is int.
- The unit is rows.

#### Application Rows Inserted

- The number of row insertions attempted. The value format is an integer. Use this attribute to gain insight into the current level of activity within the database manager.
- The type is int.
- The unit is inserts.

#### Application Rollback SQL Stmts

- The total number of SQL ROLLBACK statements that have been attempted. The value format is an integer. A rollback can result from an application request, a deadlock, or an error situation. This attribute counts only the number of rollback statements issued from applications.
- The type is int.
- The unit is rollbacks.

#### Application Total Hash Loops

- The total number of times that a single partition of a hash join was larger than the available sort heap space. The value format is an integer. Values for this attribute indicate inefficient execution of hash joins. This might indicate that the sort heap size is too small or the sort heap threshold is too small.
- The type is int.
- The unit is occurrences.

#### Application Open Rem Curs Blk

- The number of remote blocking cursors currently open for this application. The value format is an integer. Use this attribute with the Open Remote Cursors attribute to calculate the percentage of remote cursors that are blocking cursors.
- The type is int.
- The unit is cursors.

#### Application Pool Write Time

- The total amount of time spent physically writing data or index pages from the buffer pool to disk. The value format is an integer. Use this attribute with the Buffer Pool Data Writes and Buffer Pool Index Writes attributes to calculate the average page-write time. This average is important because it might indicate the presence of an I/O wait, which in turn might indicate that you must move data to a different device.
- The type is int.
- The unit is ?.

#### Application Pool Index from Estore

- Number of buffer pool index pages copied from extended storage. The value format is an integer.
- The type is int.
- The unit is pages.

#### Application Internal Commits

- The total number of commits initiated internally by the database manager. The value format is an integer.
- The type is int.
- The unit is commits.

#### Application UID SQL Stmt

- The number of SQL UPDATE, INSERT, and DELETE statements that were issued. The value format is an integer.
- The type is int.
- The unit is statements.

#### Application Direct Reads

- The number of read operations that do not use the buffer pool. The value format is an integer.
- The type is int.
- The unit is reads.

#### Application Open Rem Curs

- The number of remote cursors currently open for this application, including the cursors counted by the Open Remote Cursors with Blocking attribute. The value format is an integer.
- The type is int.
- The unit is cursors.

#### Application Rows Updated

- The number of row updates attempted. The value format is an integer. Use this attribute to gain insight into the current level of activity within the database manager.
- The type is int.
- The unit is updates.

#### Application Query Card Estimate

- An estimate of the number of rows that are returned by a query. The value format is an integer. You can compare this estimate by the SQL compiler with the actual runtime values.
- The type is int.
- The unit is rows.

#### Application Dynamic SQL Stmts

- The number of dynamic SQL statements that were attempted. The value format is an integer.
- The type is int.
- The unit is statements.

#### Application Acc Curs Blk

- The number of times that a request for an I/O block was accepted. The value format is an integer. Use this attribute with the Rejected Block Cursor Requests attribute to calculate the percentage of blocking requests that are accepted or rejected.
- The type is int.
- The unit is accepteds.

#### Application Pool Data Writes

- The number of times a buffer pool data page was physically written to disk. The value format is an integer.
- The type is int.
- The unit is writes.

#### Application Group UOW Log Space Used

- The amount of log space (in bytes) used in the current unit of work of the monitored application. Use this attribute to understand the logging requirements at the unit-of-work level. Values that are greater than or equal to 2147483647 are indicated in the portal with the Value Exceeds Maximum text, and values that are smaller than -2147483648 are indicated with the Value Exceeds Minimum text.
- The type is int.
- The unit is bytes.

#### Application Failed SQL Stmts Percent

- The percentage of SQL statements that failed to run successfully. The value format is an integer. This value is derived by dividing the value of the Failed SQL Statements attribute by the value of the Total SQL Statements attribute.
- The type is double.
- The unit is percent.

#### Application Lock Waits

- The total number of times the database applications waited for locks. The value format is an integer. At the database level, the lock waits value is the total number of times that applications waited for locks within this database. At the application-connection level, the lock waits value is the total number of times that this connection requested a lock but waited because another connection was already holding a lock on the data.
- The type is int.
- The unit is occurrences.

#### Application Total Hash Joins

- The total number of hash joins that ran. The value format is an integer.
- The type is int.
- The unit is joins.

#### Application Direct Write Reqs

- The number of requests to perform a direct write of one or more sectors of data. The value format is an integer.
- The type is int.
- The unit is requests.

#### Application Lock Wait Time

- The total elapsed time (in milliseconds) that was spent waiting for a lock.
- The type is int.
- The unit is milliseconds.

#### Application Total SQL Stmt

- The total number of dynamic and static SQL statements. This value is derived by adding the values of the Dynamic SQL Statements and the Static SQL Statements attributes.
- The type is int.
- The unit is statements.

#### Application Sort Overflows Percent

- The percentage of sorts that ran out of sort heap space and might have required disk space for temporary storage. The value format is an integer. This percentage is calculated by dividing the value of the Sort Overflows attribute by the value of the Total Sorts attribute. at the database or application level, use this attribute to evaluate the percentage of sorts that required overflow to disk. If this percentage is high, you might want to adjust the database configuration by increasing the value of the SORTHEAP configuration parameter.
- The type is double.
- The unit is percent.

#### Application Degree Parallelism

- The degree of parallelism requested when the query was bound. The value format is an integer. Use with the Agents Top attribute to determine if the query achieved maximum level of parallelism.
- The type is int.
- The unit is degree.

#### Application Total Sort Time

- The total elapsed time (in milliseconds) for all sorts that ran. The value format is an integer. at the database or application level, use this element with the Total Sorts attribute to calculate the average sort time. This average can indicate whether sorting is a performance concern.
- The type is int.
- The unit is milliseconds.

#### Application DDL SQL Stmts

- The number of SQL Data Definition Language (DDL) statements that were issued. The value format is an integer.
- The type is int.
- The unit is statements.

### **Component: Network**

Provides network information of the monitored DB2 instance.

#### **Dimensions**

##### Network IP Protocol

- The IP protocol type of the DB2 server.
- The type is int. This is a key dimension.

##### Network Listener Port

- The TCP/IP port that the database server uses in communication with a remote client.
- The type is int. This is a key dimension.

##### Network Node Name

- The format is instanceid:hostname:UD for all operating systems.
- The type is string.

##### Network DB2 Server Name

- The name of the DB2 server.
- The type is string. This is a key dimension.

##### Network IP Address

- The IP address that is used by the DB2 server.
- The type is string. This is a key dimension.

##### Network DB Partition

- The DB2 database partition node number, which can range from 0 to 999. The Aggregated and Current Partition values can be used within a query or situation filter. If you do not specify a db partition filter, data is returned for either the current database partition (single partition environment) or the aggregated database partitions (multiple partition environment). If a db partition filter is set to Aggregated, only aggregated partition data is returned. Historical data

collection includes both aggregated and individual partition attribute data. In addition to numeric partition numbers in the 0 to 999 range,.

- The type is string. This is a key dimension.

## mySAP agent metrics

The metrics for mySAP agent resource types collect data for monitoring with IBM Cloud App Management. Every mySAP agent resource type defines a set of dimensions and metrics. The descriptions provide such information as data type, dimension key, and metric unit.

### SAP ABAP Application Server

Information about mySAP Systems.

#### Component: SAP ICM Service

This data set provides information on the services that are configured for Internet Communication Manager (ICM).

##### Dimensions

IcmServices System Name

- The SAP System Identifier (SID) for the SAP system you are monitoring. For example, PRD.
- The type is string.

ICM Service Name or Port Number

- Port number or service name on which the ICMAN request accepts the corresponding protocol.
- The type is string. This is a key dimension.

Service Status

- The status of the service and whether it is currently active. The ICM does not accept requests on an inactive port.
- The type is string.

Host Name

- The fully qualified host name to which the port is linked. The valid format is S ,32.
- The type is string.

Time Period for Keep Alive ( Sec )

- Keep alive timeout in seconds. If no data is exchanged on an existing connection for this period of time, the network connection is terminated.
- The type is int.

SSL Client Verification

- The SSL Client Verification number.
- The type is int.

Virtual Host Index

- The index of the virtual host.
- The type is int.

Maximum Processing Time in Back End ( Sec )

- Timeout in seconds for processing in the backend.



- The type is int.

#### IcmServices System Label

- System label that is generated from SID\_DBhostname where SID is the target SAP system ID and DBhostname is the host name of the data base server associated with the target SAP system.
- The type is string.

#### IcmServices SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### IcmServices Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### Internet Protocol ID

- The internet Protocol ID that is used. ICM currently supports HTTP, HTTPS and SMTP.
- The type is string.

### **Component: Logon Groups**

Logon Groups is a system level data set that provides information about the logon groups and server groups used to connect users to the instances in the mySAP system.

#### **Dimensions**

##### LogonGroups System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

##### Logon/Server Group Name

- The name of the Logon/Server group that is assigned to a number of instances. Users are automatically logged on to the instance with the best response time. This attribute provides single-byte character support only. For example, ALL SERVERS is the name of the Server group you are monitoring.
- The type is string.

##### Name (Unicode)

- The name of the Logon/Server group that is assigned to a number of instances. Users are automatically logged on to the instance with the best response time. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string. This is a key dimension.

##### LogonGroups Sample Time

- The time stamp for the date and time the agent collected the data from mySAP.
- The type is timestamp.

##### sapLogonGroups Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.

- The type is string.

#### Logon Groups System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

#### Logon Groups Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### Alternate IP Address

- The alternate IP address for this instance. For example, 10. 21. 1. 11 is the name of the alternate IP address.
- The type is string.

#### Logongroups SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### Logon Groups Type

- The type of group that is being monitored.
- The type is string.

#### Logon Groups Status

- The current instance status for this Logon/Server group.
- The type is string.

#### Alternate IP Address (v4/v6)

- The alternate IP address for this instance. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

#### Statistics Sample Time

- The time stamp for the date and time that the agent created these current statistics.
- The type is timestamp.

#### Current Favorite

- The current favorite status for this instance in this logon group, which means this instance is picked for the next user that requests this logon group. For example, YES indicates that this instance is picked for the next user that requests this Logon group.
- The type is string.

#### sapLogonGroups Instance Name

- The name of the mySAP instance that is a member of this Logon/Server group. For example, ddrum2 PRD 00 is the name of the mySAP instance you are monitoring.
- The type is string.

### **Metrics**

#### Event Frequency (per/min)

- The number of events per minute on this instance. For example, 13 is the number of events per minute on this instance.
- The type is int.
- The unit is events per minute.

#### Maximum Users

- The maximum allowed number of users in this Logon group on this instance. For example, 52 is the maximum allowed number of users in this Logon group on this instance.
- The type is int.
- The unit is users.

#### Current Users

- The current number of users on this instance. For example, 9 is the current number of users on this instance.
- The type is int.
- The unit is users.

#### Current Response Time (ms)

- The current response time, in milliseconds, for this instance. For example, 56 is the number of milliseconds it takes for responses.
- The type is timestamp.
- The unit is unspecified.

#### Maximum Response Time (ms)

- The maximum allowed response time, in milliseconds, for this instance in this Logon group. For example, 0 is the maximum allowed response time for this instance in this group.
- The type is int.
- The unit is milliseconds.

### **Component: SAP ABAP Application Server Details**

SAP ABAP Application Server Details.

#### **Dimensions**

##### ABAPApplServer System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

##### Instance Host IP Address (v4/v6)

- The IP address of the physical system on which the application instance resides. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

##### Central Instance Name

- The name of the central instance application server that is configured for this mySAP system.
- The type is string.

##### ABAPApplServer Sample Time

- The time stamp for the date and time the agent collected the data from mySAP.

- The type is timestamp.

#### Message Service Configured

- A Yes/No switch to indicate if the message server is configured.
- The type is string.

#### System Description

- A user-provided description of this application instance as defined in the mySAP system transport table. This attribute provides single-byte character support only.
- The type is string. This is a key dimension.

#### Instance Stop Time

- The time stamp for the date and time the application instance stopped.
- The type is timestamp. This is a key dimension.

#### Spool Service Configured

- A Yes/No switch to indicate if the spool service is configured.
- The type is string.

#### Batch Service Configured

- A Yes/No switch to indicate if the batch service is configured.
- The type is string.

#### ABAPApplServer Description

- The dummy field for the Description column in portrait mode.
- The type is string.

#### ABAPApplServer System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### Instance Host Name

- The name of the physical system, without the domain, on which this application server resides. For example, Insthost is the name of the application instance you are monitoring.
- The type is string. This is a key dimension.

#### Instance Status

- The status of this application instance, either running or not running.
- The type is string. This is a key dimension.

#### ABAPApplServer Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### Central Instance

- A Yes/No switch to indicate if the application server is the central instance. This attribute can be useful when tailoring a situation.

- The type is string.

#### Gateway Service Configured

- A Yes/No switch to indicate if the gateway service is configured.
- The type is string.

#### ABAPApplServer SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### ABAPApplServer System Number

- The number assigned to this application server instance. For example, 01 is the number of the mySAP instance you are monitoring.
- The type is string. This is a key dimension.

#### System Description (Unicode)

- A user-provided description of this application server instance as defined in the mySAP system transport table. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### ABAPApplServer Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string. This is a key dimension.

#### Enqueue Service Configured

- A Yes/No switch to indicate if the enqueue service is configured.
- The type is string.

#### Update Service Configured

- A Yes/No switch to indicate if the update service is configured.
- The type is string.

#### System Start Time

- The time stamp for the date and time the system started.
- The type is timestamp.

#### Instance Op Mode State

- The state in which the instance is included in the current operation mode of this application server.
- The type is string.

#### Database Host IP Address

- The IP address of the physical system on which the database instance resides. This value is the same for all instances of a mySAP system. For example, 170. 106. 1. 1 is the IP address for the database host in the mySAP system you are monitoring.
- The type is string. This is a key dimension.

#### Database Name

- The name of the database instance defined for this mySAP system. This name is frequently the same as the mySAP SID, and is the same for each instance of a mySAP system. For example, DB4 is the name of the physical system on which the database server resides in the mySAP system you are monitoring.
- The type is string.

#### sapSysDetails Database Host Name

- The name of the host computer running the database instance of a system. For example, DBhost is the name of the database host in the mySAP system you are monitoring.
- The type is string. This is a key dimension.

#### Dialog Service Configured

- A Yes/No switch to indicate if the dialog service is configured.
- The type is string.

#### Operation Mode

- A text string identifier or name for the current operation mode of the system. For example, Private indicates the current operation mode of the system. This attribute provides single-byte character support only.
- The type is string.

#### Database Port

- Database Port.
- The type is int. This is a key dimension.

#### Instance Start Time

- The time stamp for the date and time the application instance started.
- The type is timestamp. This is a key dimension.

#### Instance Name

- The name of the application server.
- The type is string. This is a key dimension.

#### Instance Host IP Address

- The IP address of the physical system on which the application instance resides. For example, 170. 106. 1. 11 is the IP address of the physical system on which the application instance you are monitoring resides.
- The type is string. This is a key dimension.

#### Configuration String

- The services mask, or string, for this application server. For example, DVEBMGS indicates that the following mySAP services are configured for this instance: D = Dialog V = Update (stands for Verbucher in German) E = Enqueue B = Background M = Message server G = SNA gateway S = Spool.
- The type is string.

#### Assigned Update Instance

- The name of the application server assigned to a specific update server. For example, Updinst\_SY1\_00 is the instance configured with the mySAP update service for this application instance.

- The type is string.

#### Database Type

- Type of database.
- The type is string. This is a key dimension.

#### System Release

- The release number for the level of software installed on this application server. For example, 640 indicates the level of software installed in the SAP mySAP system you are monitoring.
- The type is string. This is a key dimension.

#### ABAPApplServer Value

- The dummy field for the Value column in portrait mode.
- The type is string.

#### Database Release

- Release associated with the database.
- The type is string. This is a key dimension.

#### Operation Mode (Unicode)

- A text string identifier or name for the current operation mode of the system. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### Database Host IP Address (v4/v6)

- The IP address of the physical system on which the database instance resides. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

#### Update2 Service Configured

- A Yes/No switch to indicate if the Update2 service is configured.
- The type is string.

### **Metrics**

#### Update2 Stopped Percent

- Percent of Update2 work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### Update2 Queue Percent

- Percentage of the dispatcher queue allotted for an Update2 that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### Spool Running Percent

- Percent of Spool work processes in the Running state.
- The type is int.

- The unit is percent.

#### Dialog Queue

- The number of tasks in the dispatch queue waiting for a Dialog work process.
- The type is int.
- The unit is tasks.

#### Spool Queue Percent

- Percentage of the dispatcher queue allotted for Spool that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### Dialog Processes

- The number of dialog processes running on this application instance.
- The type is int.
- The unit is processes.

#### Update Running Percent

- Percent of Update work processes in the Running state.
- The type is int.
- The unit is percent.

#### Update Complete Percent

- Percent of Update work processes in the Complete state.
- The type is int.
- The unit is percent.

#### Spool Processes

- The number of spool processes running on this application instance.
- The type is int.
- The unit is spool processes.

#### Update2 Processes

- Number of Update2 work processes running on this application instance.
- The type is int.
- The unit is processes.

#### Dialog Stopped Percent

- Percent of Dialog work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### Instance Up Duration

- The amount of time, in minutes, an application instance has been up in this system. For example, 12 indicates that a particular instance has been up for 12 minutes. A value of -1 indicates that there is no data at this time.
- The type is int.



- The unit is minutes.

#### System Up Duration

- The amount of time, in minutes, that the system has been up. For example, 12 indicates that the system has been up for 12 minutes. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### Enqueue Waiting Percent

- Percent of Enqueue work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### Total Active Users

- Number of active users currently for this server. It includes RFC users and interactive users.
- The type is int.
- The unit is active usres.

#### Registered Users

- Number of total registered users currently for this SAP system.
- The type is int.
- The unit is registered users.

#### Instances ConnectionFailed

- The total number of instances that have lost connection in the system.
- The type is int.
- The unit is instances.

#### Update2 Running Percent

- Percent of Update2 work processes in the Running state.
- The type is int.
- The unit is percent.

#### Interactive Users

- Number of interactive (GUI) users currently for this server.
- The type is int.
- The unit is interactive users.

#### Dialog Complete Percent

- Percent of Dialog work processes in the Complete state.
- The type is int.
- The unit is percent.

#### Enqueue Stopped Percent

- Percent of Enqueue work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### Batch Job Queue

- The number of batch jobs in Ready state.
- The type is int.
- The unit is jobs.

#### Total RFC Sessions

- The total number of RFC sessions.
- The type is int.
- The unit is sessions.

#### Instances Down

- The total number of application instances that are down in this system. This values are only reported for instances defined in an operation mode profile. For example, 3 indicates that 3 of the instances you are monitoring are not running.
- The type is int.
- The unit is instances.

#### Dialog Queue Percent

- Percentage of the dispatcher queue allotted for Dialog that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### Batch Processes

- The number of batch processes running on this application instance.
- The type is int.
- The unit is processes.

#### Update Waiting Percent

- Percent of Update work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### Batch Waiting Percent

- Percent of Batch work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### Batch Running Percent

- Percent of Batch work processes in the Running state.
- The type is int.
- The unit is percent.

#### Enqueue Running Percent

- Percent of Enqueue work processes in the Running state.
- The type is int.
- The unit is percent.

#### Update Processes

- The number of update processes running on this application instance.
- The type is int.
- The unit is processes.

#### Enqueue Complete Percent

- Percent of Enqueue work processes in the Complete state.
- The type is int.
- The unit is Percent.

#### Spool Waiting Percent

- Percent of Spool work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### Update Queue Percent

- Percentage of the dispatcher queue allotted for Update that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### Total External Sessions

- The total number of user sessions (GUI and RFC).
- The type is int.
- The unit is user sessions.

#### Instances Running

- The total number of instances that are running in this system. For example, 15 indicates that 15 instances you are monitoring are running.
- The type is int.
- The unit is instances.

#### Dialog Running Percent

- Percent of Dialog work processes in the Running state.
- The type is int.
- The unit is percent.

#### Enqueue Queue

- The number of tasks in the dispatch queue waiting for an Enqueue work process.
- The type is int.
- The unit is tasks.

#### Enqueue Queue Percent

- Percentage of the dispatcher queue allotted for Enqueue that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### Instance Down Duration

- The amount of time, in minutes, an application instance has been down. For example, 12 indicates that a particular instance has been down for 12 minutes. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### Batch Complete Percent

- Percent of Batch work processes in the Complete state.
- The type is int.
- The unit is percent.

#### Spool Complete Percent

- Percent of Spool work processes in the Complete state.
- The type is int.
- The unit is percent.

#### Enqueue Processes

- Number of enqueue work processes running on this application instance.
- The type is int.
- The unit is processes.

#### Spool Queue

- The number of tasks in the dispatch queue waiting for a Spool work process.
- The type is int.
- The unit is tasks.

#### Dialog Waiting Percent

- Percent of Dialog work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### RFC Users

- Number of RFC users currently for this server.
- The type is int.
- The unit is RFC users.

#### Total GUI Sessions

- The total number of non-APPC-TM GUI sessions.
- The type is int.
- The unit is sessions.

#### Instances Passive

- The total number of instances that are in passive state in the system.
- The type is int.
- The unit is instances.

#### Update Stopped Percent

- Percent of Update work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### Batch Stopped Percent

- Percent of Batch work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### Update2 Queue

- The number of tasks in the dispatch queue waiting for an Update2 work process.
- The type is int.
- The unit is tasks.

#### Update Queue

- The number of tasks in the dispatch queue waiting for an Update work process.
- The type is int.
- The unit is tasks.

#### Update2 Waiting Percent

- Percent of Update2 work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### Spool Stopped Percent

- Percent of Spool work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### Active Users

- The current number of users logged on to this application instance. For example, 47 indicates the number of users currently logged on to the instance you are monitoring.
- The type is int.
- The unit is users.

#### Update2 Complete Percent

- Percent of Update2 work processes in the Complete state.
- The type is int.
- The unit is percent.

#### NoWP Queue

- The number of tasks in the dispatch queue waiting to be processed by the dispatcher itself or some other system service.
- The type is int.
- The unit is tasks.

## Component: SAP HTTP Services

This data set shows HTTP Services details.

### Dimensions

#### Client

- The client that is used to connect to the SAP server.
- The type is string.

#### Changed On Timestamp

- The time and date when the HTTP services were changed.
- The type is timestamp.

#### Created on Timestamp

- The date and time on which the HTTP services were created.
- The type is timestamp.

#### Session Timeout

- Session Timeout for a stateful connection in time format.
- The type is string.

#### HttpServices Managed System

- The identifier for this SAP resource.
- The type is string. This is a key dimension.

#### Created for Client

- The Internet Communication Framework (ICF) created for the client.
- The type is string.

#### HttpServices System Name

- The SAP System Identifier (SID) for the SAP system that you are monitoring. For example, PRD.
- The type is string.

#### Changed for Client

- Client for HTTP service changed.
- The type is string.

#### Status

- Status of the service, for example, Active or Inactive.
- The type is string.

#### sapHttpServices SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### sapHttpServices Changed On

- The date on which the HTTP services were changed.
- The type is string.

#### Parent GUID

- GUID of the parent node.
- The type is string. This is a key dimension.

#### Service Name

- Name of a Service in Internet Communication Framework.
- The type is string. This is a key dimension.

#### Created on

- The date on which the HTTP services were created.
- The type is string.

#### saphttpServices Host Number

- Number of a Virtual Host.
- The type is int.

#### sapHttpServices System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the data base server associated with the target SAP system.
- The type is string.

#### Path

- Path of the HTTP service.
- The type is string.

#### SAP Authority

- Authorization to use an ICF service.
- The type is string.

#### Session Timeout(Sec.)

- Session Timeout for a stateful connection in seconds.
- The type is int.

#### Service Node GUID

- GUID of the ICF Service node.
- The type is string.

#### User

- Logon name of the user.
- The type is string.

#### sapHttpServices Description

- Description of the HTTP Service.
- The type is string.

#### sapHttpServices Host Name

- Host of the service.
- The type is string.

#### sapHttpServices Created By

- The HTTP Services that are created by default and also those HTTP services that are created by the user.
- The type is string.

#### Last Changed By

- The user name of the person who last changed the HTTP Service.
- The type is string.

### **Component: Output Requests**

Output Requests is a system level data set that provides information about all output requests in the mySAP system.

#### **Dimensions**

##### OutputReq Sample Interval End

- The time stamp for the stopping time of the data supplied by the SAP agent.
- The type is timestamp.

##### Output Request Client

- A text string identifier or name for the originating client. For example, A800 indicates the identifier for the originating client.
- The type is string.

##### OutputReq System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

##### OutputReq Sample Interval Start

- The time stamp for the beginning time of the data supplied by the SAP agent.
- The type is timestamp.

##### sapOutputRequests Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

##### sapOutputRequests Output Device (Unicode)

- A text string identifier or name for the current output device. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string. This is a key dimension.

##### Output Requests Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

##### Spooler System Name

- A text string identifier or name for the system where the spooler is running. For example, DDRUM2\_PRD indicates the system where the spooler is running.



- The type is string.

#### Department (Unicode)

- A text string identifier or name for the current department receiving the request. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### Output Requests Recipient

- A text string identifier or name for the current recipient of the request. For example, RBROWN indicates the name of the recipient for the request.
- The type is string.

#### sapOutputReq Output Format

- A text string identifier for the current output format. This attribute provides single-byte character support only. For example, X\_65\_255 indicates the output format.
- The type is string.

#### Print Request Time

- The time stamp for the date and time the print request was created.
- The type is timestamp.

#### Output Requests Creator

- The user ID for the originator of the request. For example, RSMITH indicates the originator of the request.
- The type is string.

#### Spooler Host Name

- A text string identifier or name for the host where the spooler is running. For example, DDRUM2 indicates the name of the spooler host.
- The type is string.

#### sapOutputReq Output Format (Unicode)

- A text string identifier for the current output format. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### Output Requests System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

#### OutputReq Spool Title

- A text identifier or name for the spool file. This attribute provides single-byte character support only. For example, PRINTA indicates the title of the spool file.
- The type is string.

#### OutputReq Spool Number

- A numeric identifier for the spool file. For example, 31806 indicates the numeric identifier for the spool file.

- The type is int.

#### Print Reason

- The reason for the print request.
- The type is string.

#### Output Requests Department

- A text string identifier or name for the current department receiving the request. This attribute provides single-byte character support only. For example, DEV indicates the current department receiving the request.
- The type is string.

#### Spool Title (Unicode)

- A text identifier or name for the spool file. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### Host Spool Id

- A text string identifier for the print host spooler. For example, CAN2 indicates the identifier for the host spooler.
- The type is string.

#### sapOutputRequests Output Device

- A text string identifier or name for the current output device. This attribute provides single-byte character support only. For example, LP01 indicates the name of the output device.
- The type is string.

#### Print Status

- The status of a print request.
- The type is string.

#### OutputReq SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

### **Metrics**

#### Error Print Requests

- The number of print requests with errors. For example, 3 indicates the number of print requests with errors.
- The type is int.
- The unit is requests.

#### sapOutputReq Processed Print Requests

- The number of processed print requests. For example, 2 indicates the number of processed print requests.
- The type is int.
- The unit is requests.

#### sapOutputRequests Number of Copies

- The number of copies requested. For example, 31 indicates the number of copies requested.
- The type is int.
- The unit is copies.

#### Print Pending (mins) (Superseded)

- Time calculated for the pending output request. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### Output Requests Size

- The amount of disk space or memory to which the request can spool. For example, 4056 indicates that 4 MB of disk space is available for the request.
- The type is int.
- The unit is MB.

#### Failed Print Requests

- The number of print requests that did not complete. For example, 2 indicates the number of print requests that did not complete.
- The type is int.
- The unit is requests.

#### Print Pending (mins)

- Time calculated for the pending output request.
- The type is int.
- The unit is minutes.

### **Component: DB2 Performance History**

Database performance history is system level data set that provides information about DB2 database performance history occurring in a SAP system.

#### **Dimensions**

##### Index Logical Reads

- Number of Index Logical Reads.
- The type is int.

##### Index Physical Writes

- Number of Index Physical Writes.
- The type is int.

##### Deadlocks

- The total number of deadlocks that have occurred since the first database connection.
- The type is int.

##### Data Physical Writes

- Number of data Physical Writes.
- The type is int.

#### Average Physical Read Time(ms)

- Average Physical Read Time in milliseconds.
- The type is int.

#### Workload

- Type of workload.
- The type is string.

#### Average Physical Write Time(ms)

- Average Physical Write Time in milliseconds.
- The type is int.

#### Lock Waits

- The total amount of time that applications or connections waited for locks.
- The type is int.

#### Db2PerformanceHistory System Name

- System Name.
- The type is string.

#### Data Physical Reads

- Number of data Physical Reads.
- The type is int.

#### Rollback Statements

- The total number of SQL ROLLBACK statements that have been attempted.
- The type is int.

#### Data Logical Reads

- Number of data Logical Reads.
- The type is int.

#### Commit Statements

- The total number of SQL COMMIT statements that have been attempted.
- The type is int.

#### Lock Wait Time (ms)

- The total elapsed time waited for a lock. Elapsed time is given in milliseconds.
- The type is int.

#### Db2PerfHistory System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the data base server associated with the target SAP system.
- The type is string.

#### Exclusive Lock Escalations

- The number of times that locks have been escalated from several row locks to one exclusive table lock, or the number of times an exclusive lock on a row caused the table lock to become an exclusive lock.
- The type is int.

#### Lock Escalations

- The number of times that locks have been escalated from several row locks to a table lock.
- The type is int.

#### Db2PerformanceHistory Managed System

- The identifier for this SAP resource.
- The type is string. This is a key dimension.

#### Db2PerformanceHistory SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### Row Insert Timestamp

- Row Insert Date and TimeThe date and time when the row is inserted.
- The type is timestamp.

#### Index Physical Reads

- Number of Index Physical Reads.
- The type is int.

### Metrics

#### Db2PerfHistory Sample Time

- The time stamp for the date and time when the agent collected data from SAP.
- The type is timestamp.
- The unit is unspecified.

### Component: Updates Information

Updates Information is a system level data set that provides information about updates to the database in the mySAP system.

#### Dimensions

##### Update Server

- The name of the server being used for record updates. For example, ddrum2\_PRD\_00 is the identifier for the server you are using.
- The type is string.

##### Program (Unicode)

- A text string identifier or name for the program that is performing the update. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### Updates Information Program

- A text string identifier or name for the program that is performing the update. This attribute provides single-byte character support only. For example, SAPLY210 identifies the name of the program associated with this process.
- The type is string.

#### Updates Information Status

- The current status of the update.
- The type is string.

#### sapUpdateInformation Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### Updates Information Time

- The time stamp for the time the update was attempted.
- The type is timestamp. This is a key dimension.

#### Status Number

- The identifier for the status. For example, 9 indicates the number of the update status.
- The type is int.

#### Error

- The type of error that occurred during an update. This attribute provides single-byte character support only. For example, 00671ABAP/4 processor POSTING\_ILLEGAL\_STATEMENT indicates that an error occurred during the execution of a particular mySAP process.
- The type is string.

#### State Description

- The description of the current state of the update. For example, Update is active indicates that an update is occurring. This attribute provides single-byte character support only.
- The type is string.

#### State Description (Unicode)

- The description of the current state of the update. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### UpdateInfo Sample Interval End

- The time stamp for the stopping time of the data supplied by the SAP agent.
- The type is timestamp.

#### Updates Information System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

#### UpdateInfo SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### State Code

- The current state of the update.
- The type is string.

#### updateInfo Sample Interval Start

- The time stamp for the beginning time of the data supplied by the SAP agent.
- The type is timestamp.

#### Status Description

- Text describing the status of the update request. For example, Update is active indicates that the current update request is active. This attribute provides single-byte character support only.
- The type is string.

#### Function Module

- The name of the function module associated with the update. For example, 03 indicates the name of the function module being used.
- The type is string.

#### Updates Information Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### Updates Information User ID

- The name of the user performing the transaction. For example, RBROWN is the name of the user performing the transaction.
- The type is string.

#### Updates Information Client

- A text string identifier or name, for the source client session. For example, 017 identifies the name of the client for this session.
- The type is string.

#### UpdateInfo System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### Status Description (Unicode)

- Text describing the status of the update request. For example, Update is active indicates that the current update request is active. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### Error (Unicode)

- The type of error that occurred during an update. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### Update Key

- The identifier for the update key. For example, 19991102131415000ddrum2... 002 is the identifier of the update key.
- The type is string. This is a key dimension.

#### UpdateInfo Transaction Code

- The identifier for the transaction code. For example, FB01 is the identifier for the transaction code you are using.
- The type is string.

### **Component: SAP SYS Connection Monitoring**

Describes Connection Monitoring.

#### **Dimensions**

##### connectionMonitoring System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the database server associated with the target SAP system.
- The type is string.

##### connectionMonitoring Last Changed By

- Name of the user who last changed the RFC destination.
- The type is string.

##### connectionMonitoring SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

##### connectionMonitoring Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

##### RFC Status

- The number of Inactive RFC Connections in the system.
- The type is int.

##### RFC Type

- Type of the RFC destination that is being monitored.
- The type is string.

##### Description

- Description of the RFC connection.
- The type is string.

##### Created On

- The time when the RFC destination was created.
- The type is timestamp.

##### Changed On



- The time when the RFC destination was last changed.
- The type is timestamp.

#### RFC Connection

- Name of the RFC destination that is specified in the function call.
- The type is string.

#### connectionMonitoring System Name

- System ID of the SAP System running on the server.
- The type is string. This is a key dimension.

#### Created By

- Name of the user who created the RFC destination.
- The type is string.

### Metrics

#### Count RFC Status

- The total count of the RFC connections. It is used to plot the graph in the Connection Monitoring workspace.
- The type is int.
- The unit is connections.

#### connectionMonitoring Sample Time

- The time stamp for the date and time the agent collected data from mySAP system.
- The type is timestamp.
- The unit is unspecified.

### Component: SAP Batch Data create

Batch Data Create is a system level data set that provides information about the configuration, progress, and performance of BDC sessions in the mySAP system.

#### Dimensions

##### Session Name (Unicode)

- A text string identifier or name for the BDC session. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string. This is a key dimension.

##### Batch Data Create Creator

- A text string identifier or user ID for the user who created the session. For example, RSMITH indicates the name of the person who created the session.
- The type is string. This is a key dimension.

##### Queue Id

- The BDC queue Id from APQI-QID.
- The type is string.

##### BatchDataCReate Sample Interval End

- The time stamp for the specific date and time that the collection period stopped.
- The type is timestamp.

#### Batch Data Create Created

- The time stamp for the date and time the BDC session or range of sessions occurred.
- The type is timestamp.

#### Batch Data Create Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### Start Mode

- The process used to begin the session.
- The type is string.

#### Batch Data Create Authorization

- ID for the permission for the session, a text string. For example, RSMITH indicates the person who authorized the session.
- The type is string.

#### Session Name

- A text string identifier or name for the BDC session. This attribute provides single-byte character support only. For example, RSMITH081358 indicates the name of the session.
- The type is string.

#### sapBatchDataCreate Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### BatchDataCreate System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the database server associated with the target mySAP system.
- The type is string.

#### BatchDataCreate Sample Interval Start

- The time stamp for the specific date and time that the collection period started.
- The type is timestamp.

#### Batch Data Create System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

#### sapBatchDataCreate SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### Last Changed

- The time stamp for the date and time the session was most recently modified.

- The type is timestamp.

#### Locked Until

- The time stamp for the specific date and time, or range, before which this session or a range of sessions cannot be processed.
- The type is timestamp.

#### Batch Data Create Status

- The status for the session.
- The type is string.

#### Batch Data Create Client

- The name of the client session. For example, 800 identifies the source client for this session.
- The type is string.

### **Metrics**

#### Pending Transactions

- The number of transactions not yet completed in this BDC session. For example, 4 identifies the number of transactions not yet completed in this session.
- The type is int.
- The unit is transaction.

#### Error Screens

- The number of screens with errors. For example, 2 indicates the number of screens with errors.
- The type is int.
- The unit is screens.

#### Total Transactions

- The total number of transactions in this BDC session. For example, 67 identifies the number of transactions for this session.
- The type is int.
- The unit is transaction.

#### Error Transactions

- The number of transactions with errors. For example, 4 indicates the number of transactions with errors.
- The type is int.
- The unit is transactions.

#### Total Screens

- The total number of screens in this BDC session. For example, 6 indicates the total number of screens for this session.
- The type is int.
- The unit is screens.

#### Pending Screens

- The number of screens not yet completed in this BDC session. For example, 2 indicates the number of screens not yet completed.

- The type is int.
- The unit is screens.

#### Completed Screens

- The number of completed screens in this BDC session. For example, 21 indicates the number of screens that completed.
- The type is int.
- The unit is Screens.

#### Completed Transactions

- The number of completed transactions in this BDC session. For example, 35 indicates the number of transactions that completed.
- The type is int.
- The unit is transactions.

#### Deleted Screens

- The number of deleted screens in this BDC session. For example, 3 indicates the number of screens that were deleted in this session.
- The type is int.
- The unit is screens.

#### Deleted Transactions

- The number of deleted transactions in this BDC session. For example, 4 indicates the number of transactions that were deleted in this session.
- The type is int.
- The unit is transactions.

### **Component: SAP SYS ABAP Dump Details**

Provides information about ABAP short dumps occurring in the mySAP system.

#### **Dimensions**

##### Line Number

- The numeric identifier for the line of code in the ABAP INCLUDE where the error occurred.
- The type is int.

##### Mode Number

- The mode number for the ABAP dump.
- The type is string. This is a key dimension.

##### dumpDetails Create Time

- The time stamp for the date and time the ABAP dump was created.
- The type is timestamp. This is a key dimension.

##### Hold Status

- The hold status for the ABAP dump. One of the following values is possible: X = Held F = Free.
- The type is string.

##### dumpDetails System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### Line Number (Superseded)

- The numeric identifier for the line of code in the ABAP INCLUDE where the error occurred. For example, 750 indicates the line in the code where the error occurred.
- The type is int.

#### dumpDetails SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### Dump Title

- A text string identifier or name for the ABAP dump that was created. For example, DBIF RSQL INVALID CURSOR indicates the name of the ABAP dump.
- The type is string.

#### dumpDetails User ID

- The text string identifier for the person who generated the ABAP dump. For example, LSMITH is the name of the person who generated the ABAP dump.
- The type is string. This is a key dimension.

#### dumpDetails Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### dumpDetails Host

- A text string identifier or name for the computer serving as the host where the ABAP dump was created. For example, ddrum2 indicates the name of the host where the ABAP dump was created.
- The type is string.

#### Include Name (Unicode)

- A text string identifier or name for the ABAP INCLUDE name. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### Program Name

- The text string identifier for the ABAP program that generated the ABAP dump. This attribute provides single-byte character support only. For example, SAPLY210 indicates the name of the ABAP program.
- The type is string.

#### Program Name (Unicode)

- The text string identifier for the ABAP program that generated the ABAP dump. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### Include Name

- A text string identifier or name for the ABAP INCLUDE name. This attribute provides single-byte character support only. For example, LY210U58 indicates the ABAP INCLUDE name.
- The type is string.

#### dumpDetails Sample Interval End

- The time stamp for the stopping time of the data supplied by the SAP agent.
- The type is timestamp.

#### dumpDetails System Name

- Not Available.
- The type is string. This is a key dimension.

#### dumpDetails Sample Interval Start

- The time stamp for the beginning time of the data supplied by the SAP agent.
- The type is timestamp.

#### dumpDetails Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### dumpDetails Client

- A text string identifier or name for the source client where the ABAP dump was created.
- The type is string.

### **Component: SAP SYS Outbound Queues overview**

Provides information about the outbound queues, the destination status, and error messages associated with the queues.

#### **Dimensions**

##### tRFC First Counter

- First counter for the serialized tRFC.
- The type is string.

##### outboundQueues Managed System

- The identifier for this SAP resource.
- The type is string. This is a key dimension.

##### outboundQueues Queue Version

- Current qRFC version.
- The type is string.

##### First Transaction ID

- First Transaction ID of the Logical Unit of Work (LUW) of the Outbound queue.
- The type is string.

##### First Application Server Timestamp

- First execution time stamp.

- The type is timestamp.

#### Last Application Server Timestamp

- Last execution time stamp. The valid format is time stamp.
- The type is timestamp.

#### outboundQueues Queue Name

- Name of the qRFC outbound Queue.
- The type is string. This is a key dimension.

#### outboundQueues Queue Count

- Number of queues grouped by queue status. This attribute is for internal use in query.
- The type is int.

#### outboundQueues SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### tRFC Last Counter

- Last counter for the serialized tRFC.
- The type is string.

#### outboundQueues Queue Destination

- qRFC outbound queue logical destination in function call. The valid format is an alphanumeric string, with a maximum of 32 characters.
- The type is string. This is a key dimension.

#### outboundQueues Queue Status

- Status of the qRFC outbound queue.
- The type is string.

#### Queue Error Message

- Outbound queue error message according to status of the queue. The valid format is an alphanumeric string, with a maximum of 73 characters.
- The type is string.

#### outboundQueues Client

- Client ID of the current User.
- The type is string.

#### outboundQueues System Name

- System ID of the SAP System running on server.
- The type is string. This is a key dimension.

#### outboundQueues Queue Supplement

- Current qRFC supplement number.
- The type is int.

outboundQueues System Label

- System label generated from the SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the database server associated with the target SAP system.
- The type is string.

Wait for Queue

- Name of the queue for which the current queue execution is waiting.
- The type is string.

Queue Counter In LUW

- Outbound queue Logical Unit of Work (LUW) Counter within a transaction.
- The type is string.

### **Metrics**

outboundQueues Queue Entries

- Number of queue Logical Unit of Work (LUW) entries.
- The type is int.
- The unit is LUW entries.

## **Component: SAP Message Server Monitor**

This data set shows the detailed information about the client Message Server Monitor for a given SAP System.

### **Dimensions**

MsgServerMonitor System Name

- The SAP System Identifier (SID) for the SAP system you are monitoring. For example, PRD.
- The type is string.

Field Name

- Shows Message Server Monitor Information.
- The type is string.

MsgServerMonitor SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

MsgServerMonitor Managed System

- The identifier for this SAP resource.
- The type is string. This is a key dimension.

MsgServerMonitor System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the data base server associated with the target SAP system.
- The type is string.

### **Metrics**

Field Value

- Shows the Message Server Monitor value.



- The type is string.
- The unit is unspecified.

### **Component: SAP Office Inbox**

SAP Office Inbox is a system level data set that provides information about SAP office resources and mail in the mySAP system.

#### **Dimensions**

sapSAPOfficeInbox Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

SAP Office Inbox System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

Attachment Type

- The type of mail item as specified by its file extension. For example, DOC, XLS, TXT, and so on. DOC indicates the type of mail item is a document type.
- The type is string.

User Name

- The name of the user who owns the SAP Office inbox. For example, LEROY BROWN is the name of the user who owns the SAP Office inbox.
- The type is string.

SAP Office Inbox Priority

- The priority of the mail item (the higher the number, the higher the priority). For example, 9 indicates that the mail item is of a high priority.
- The type is int.

SAP Office Inbox Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

SAP Office Inbox User ID

- The identifier for the user who owns the SAP Office inbox. For example, LGREEN is the name of the user who owns the SAP Office inbox.
- The type is string. This is a key dimension.

Received Time

- The time stamp for the date and time when the mail item was received.
- The type is timestamp.

SAP Office Inbox Owner

- The user name of the person who currently owns the mail item. For example, LGREEN indicates the name of the person who owns the mail item.
- The type is string.

#### Mail Name (Unicode)

- The name of the mail item. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### Action Name

- The name of the mySAP action specified in the mail item, such as program name, function module, or transaction name. This attribute provides single-byte character support only. For example, Y\_210\_NOTIFY indicates the name of the action in progress.
- The type is string.

#### Mail Title (Unicode)

- The title of the mail item. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### SAP Office Inbox Changeable

- An indicator of whether a mail item is modifiable.
- The type is string.

#### OfficeInbox Sample Interval Start

- The time stamp for the beginning time of the data supplied by the SAP agent.
- The type is timestamp.

#### Size (bytes)

- The size, in bytes, of the mail item. For example, 1785 indicates the size of the mail item.
- The type is int.

#### SAP Office Inbox Status

- The status of the mail item.
- The type is string.

#### Open Time

- The time stamp for the date and time when the mail item was opened and viewed.
- The type is timestamp.

#### Action Name (Unicode)

- The name of the mySAP action specified in the mail item, such as program name, function module, or transaction name. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### SAP Office Inbox Author

- The name of the user who created the mail item. For example, WBROWN indicates the name of the user.
- The type is string.

#### SOOfficeInbox ample Interval End

- The time stamp for the stopping time of the data supplied by the SAP agent.
- The type is timestamp.

#### OfficeInbox System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### Action Type

- The type of mySAP action specified in the mail item, such as program, function module name, or transaction. This attribute provides single-byte character support only. For example, FUNCTION MODULE indicates the type of SAP Office action associated with the mail item.
- The type is string.

#### SAP Office Inbox Express

- Indicator of whether the mail item is an Express mail type or not.
- The type is string.

#### Mail Name

- The name of the mail item. This attribute provides single-byte character support only. For example, NOTE indicates the name of the mail item.
- The type is string.

#### Mail Title

- The title of the mail item. This attribute provides single-byte character support only.
- The type is string.

#### SAP Office Inbox Sensitivity

- The sensitivity of the mail item.
- The type is string.

#### OfficeInboxSAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### Number of Attachments

- The number of attachments included with the mail item. For example, 3 indicates the three attachments are included with the mail item.
- The type is int.

#### Sent Time

- The time stamp for the date and time the mail item was sent.
- The type is timestamp.

#### Mail Type

- The type of mail item in the inbox. For example, Office, Workflow, or Deadline. This attribute provides single-byte character support only. Workflow indicates the mail item is of the Workflow type.

- The type is string.

#### Expiration Time

- The time stamp for the expiration date and time of the mail item.
- The type is timestamp.

#### Size (bytes) (Supperseded)

- The size, in bytes, of the mail item. For example, 1785 indicates the size of the mail item.
- The type is int.

#### Mail Type (Unicode)

- The type of mail item in the inbox. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### SAP Office Inbox Client

- A text string identifier, or number, for the execution client. For example, 800 indicates the client.
- The type is string.

#### Action Type (Unicode)

- The type of mySAP action specified in the mail item, such as program, function module name, or transaction. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

### Metrics

#### Inbox Pending Time (mins)

- The amount of time, in minutes, that the mail item spent in the inbox prior to being opened.
- The type is int.
- The unit is minutes.

#### Inbox Pending Time (mins) (Supperseded)

- The amount of time, in minutes, that the mail item spent in the inbox prior to being opened. For example, 1171 indicates that the mail item spent 1,171 minutes in the inbox before being opened.
- The type is int.
- The unit is minutes.

### Component: SAP SYS Lock Entries

Provides information about locked objects in the mySAP system.

#### Dimensions

##### lockDetails SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

##### Backup Flag

- The identifier for the backup flag. For example, Y indicates that the backup flag is set.

- The type is string.

#### Update Owner Name

- Contains the ID of the Lock User of the Logical Unit of Work (LUW)-Update Task.
- The type is string.

#### lockDetails System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

#### lockDetails Create Time

- The time stamp for the date and time the lock was created.
- The type is timestamp.

#### lockDetails Argument

- The argument value (key fields) of a lock entry. An entry locks the entries in a table that are specified by the argument value. For example, SAPLY210 is an example of an argument value.
- The type is string.

#### lockDetails Sample Time

- The time stamp for the date and time the agent collected the data from mySAP.
- The type is timestamp.

#### lockDetails Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### lockDetails System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### Owner Name

- Contains the ID of the Lock User of the Logical Unit of Work (LUW).
- The type is string.

#### SyslockDetails Transaction Code

- The identifier for the transaction code. For example, SMLG is the identifier for the transaction code.
- The type is string.

#### lockDetails Client

- A text string identifier or number for the originating client. For example, 800 indicates the client.
- The type is string.

#### Work Process

- The numeric identifier for the work process. For example, 3 is the number of the work process.
- The type is int.

#### Lock Object Name

- The name of the object being locked. For example, ES\_RZL\_LIP indicates the name of the object being locked.
- The type is string. This is a key dimension.

#### lockDetails User ID

- The name of the user who has set a lock. For example, RBROWN is the name of the user generating locks.
- The type is string.

#### lockDetails System Number

- The identifier for the mySAP system you are monitoring. For example, 06 is the name of the mySAP system you are monitoring.
- The type is string.

#### lockDetails Owner

- The name of the person associated with the lock. For example, LGREEN indicates the name of the person generating the lock.
- The type is string.

#### lockDetails Host

- A text string identifier or name for the computer serving as the host. For example, agoura1 indicates the identifier for the host.
- The type is string.

#### locksDetails Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### lockDetails Group

- The name of the group associated with the lock. For example, RZLLITAB indicates the name of the lock group.
- The type is string.

#### Update Owner

- The identifier for the person who holds the locks for update. For example, ddrum2..0002199901041 is the identifier for the person who holds the locks for update.
- The type is string.

### Metrics

#### Hold Count

- The total number of locks held. For example, 1 indicates the total number of locks held.
- The type is int.
- The unit is locks.

#### Update Hold Count

- The total number of locks held for update. For example, 2 indicates the total number of locks held for update.

- The type is int.
- The unit is locks.

Lock Age (mins)

- The amount of time, in minutes, elapsed since the lock was created.
- The type is int.
- The unit is minutes.

Lock Age (mins) (Superseded)

- The amount of time, in minutes, elapsed since the lock was created. For example, 33 indicates the number of minutes elapsed since the lock was created.
- The type is int.
- The unit is minutes.

### **Component: SAP SYS Inbound Queues Overview**

Provides information about the inbound queues, its destination status and error messages.

#### **Dimensions**

Queue Error Messages

- Inbound queue error message according to the status of the queue.
- The type is string.

inboundQueues System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the database server associated with the target SAP system.
- The type is string.

inboundQueues Client

- Client ID of the current User.
- The type is string.

First Timestamp

- First Execution time stamp.
- The type is timestamp.

Last Timestamp

- Last Execution time stamp.
- The type is timestamp.

First TID

- First transaction ID of the Logical Unit of Work (LUW) of the inbound queue.
- The type is string.

inboundQueues Queue Status

- qRFC inbound queue status.
- The type is string.

tRFC First Count

- First counter for the serialized tRFC.
- The type is string.

#### Queue LUW Counter

- Inbound queue Logical Unit of Work (LUW) counter within a transaction.
- The type is string.

#### inboundQueues Queue Version

- Current qRFC version.
- The type is string.

#### inboundQueues Managed System

- The identifier for this SAP resource.
- The type is string. This is a key dimension.

#### inboundQueues SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### inboundQueues Queue Count

- Number of queues group by queue status.
- The type is int.

#### tRFC Last Count

- Last counter for the serialized tRFC.
- The type is string.

#### inboundQueues System Name

- System ID of the SAP System running on the server.
- The type is string. This is a key dimension.

#### inboundQueues Queue Name

- Name of qRFC Inbound Queue.
- The type is string. This is a key dimension.

#### inboundQueues Queue Supplement

- Current qRFC supplement number.
- The type is int.

#### inboundQueues Queue Destination

- qRFC inbound queue logical destination in function call. The valid format is an alphanumeric string, with a maximum of 32 characters.
- The type is string. This is a key dimension.

### **Metrics**

#### inboundQueues Queue Entries

- Number of Logical Unit of Work (LUW) entries of qRFC inbound queue.
- The type is int.



- The unit is LUW entries.

## Component: DB2 Configuration Information

This data set shows details of the DB2 configuration.

### Dimensions

#### Lock List Before Escalation(%)

- This parameter defines a percentage of the lock list held by an application that must be filled before the database manager performs escalation. When the number of locks held by any one application reaches this percentage of the total lock list size, lock escalation occurs for the locks held by that application. Lock escalation also occurs if the lock list runs out of space.
- The type is float.

#### Statement Heap Size(Pages)

- Specifies the size of the statement heap that is used as a work space for the SQL or XQuery compiler during compilation of an SQL or XQuery statement. Unit of measure is Pages (4 KB).
- The type is int.

#### Database Release Level

- The release level of the database manager that uses the database. If a database upgrade does not complete or fails, this parameter shows the release level of the database before the upgrade. This release level can differ from the release parameter that is associated with the release level of the database configuration file.
- The type is int.

#### Number of Database Backups to Retain

- Specifies the number of database backups to retain for a database. After the specified number of backups is reached, old backups are marked as expired in the recovery history file.
- The type is int.

#### Number of I/O Servers

- Specifies the number of I/O servers for a database. A database can not have any more than this number of I/O servers for prefetching and utilities in progress at any time.
- The type is int.

#### Average Number of Active Applications

- Used by the query optimizer to help estimate how much buffer pool space is available at run time for the access plan chosen.
- The type is int.

#### Statistics Heap Size(Pages)

- Indicates the maximum size of the heap that is used in collecting statistics by using the RUNSTATS command. Unit of measure is Pages (4 KB).
- The type is int.

#### Catalog Cache Size(Pages)

- The maximum space in pages that the catalog cache uses from the database heap. In a partitioned database system, there is one catalog cache for each database partition. The unit of measure is Pages (4 KB).
- The type is int.

#### Number of Sorts Since First Connect

- Number of sorts since first connect.
- The type is int.

#### Db2ConInfo Managed System

- The identifier for this SAP resource.
- The type is string. This is a key dimension.

#### Deadlocks Since First DB Connect

- The total number of deadlocks that have occurred since the first database connection.
- The type is int.

#### Restore Pending

- States whether a RESTORE PENDING status exists in the database.
- The type is string.

#### Log File Size(Pages)

- Defines the size of each primary and secondary log file. The size of these log files determines and limits the number of log records that you can write to them before they become full and a new log file is required. The unit of measure is Pages (4 KB).
- The type is int.

#### Maximum Number of Active Applications

- Specifies the maximum number of concurrent applications that you connect, both local and remote, to a database.
- The type is int.

#### Number of Asynchronous Page Cleaners

- Specifies the number of asynchronous page cleaners for a database. These page cleaners write changed pages from the buffer pool to disk before the space in the buffer pool is required by a database agent.
- The type is int.

#### Locks Currently Held

- This parameter shows the number of locks currently held. If the monitor information is at the database level, this is the total number of locks currently held by all applications in the database.
- The type is int.

#### Rollforward Pending Indicator

- Informs you whether or not a roll forward recovery is required, and where it is required. The recovery (using ROLLFORWARD DATABASE) must complete before you can access the database or table space.
- The type is string.

#### Application Heap Size(Pages)

- Defines the number of private memory pages that are available for use by the database manager on behalf of a specific agent or subagent. This parameter is allocated when an agent or subagent is initialized for an application. The Unit of Measure is Pages. Each Page size is 4 KB.

- The type is int.

#### Number of Secondary Log Files

- Specifies the number of secondary log files that are created and used for recovery log files. The valid format is a 4-byte integer.
- The type is int.

#### Db2ConInfo System Name

- The SAP System Identifier (SID) for the SAP system you are monitoring. For example, PRD.
- The type is string.

#### Backup Pending Indicator

- Indicates that you must do a full backup of the database before accessing it. This parameter is on only if the database configuration is changed.
- The type is string.

#### Db2ConnInfo SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### Auto Restart

- Determines whether the database manager can, in the event of an abnormal termination of the database, automatically call the restart database utility when an application connects to a database. The default value is ON.
- The type is string.

#### Dynamic Query Management

- This parameter determines whether Query Patroller captures information about submitted queries. If this parameter is set to ENABLE, Query Patroller captures information about the query. If parameter is set to DISABLE, Query Patroller does not capture any information about submitted queries.
- The type is string.

#### Db2ConnInfo System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the data base server associated with the target SAP system.
- The type is string.

#### Sort Heap Size(Pages)

- Defines the maximum number of private memory pages used for private sorts, or the maximum number of shared memory pages used for shared sorts. Unit of measure is Pages (4 KB).
- The type is int.

#### Database Heap Size(Pages)

- The maximum memory used by the database heap. There is one database heap per database and the database manager uses it for the applications that are connected to the database. The unit of measure is Pages (4 KB).
- The type is int.

#### Territory of the Database

- Shows the territory used to create the database. Territory is used by the database manager when processing data that is sensitive to territory.
- The type is string.

#### Total Sort Heap Allocated

- The total number of allocated pages of sort heap space for all sorts at the level chosen and at the time the snapshot was taken.
- The type is int.

#### Package Cache Size(Pages)

- Allocated out of the database shared memory, and is used for caching of sections for static and dynamic SQL and XQuery statements on a database. Unit of measure is Pages (4 KB).
- The type is int.

#### Maximum Number of Database Files Open per Application

- Specifies the maximum number of file handles that you open per application.
- The type is int.

#### Lock Waits Since First Connect(microSec)

- The time taken in microseconds that applications or connections waited for locks.
- The type is int.

#### Utility Heap Size(Pages)

- Indicates the maximum amount of memory that is be used simultaneously by the BACKUP, RESTORE, and LOAD (including load recovery) utilities.
- The type is int.

#### Maximum Storage for Lock List (Pages)

- Indicates the amount of storage that is allocated to the lock list. There is one lock list per database and it contains the locks held by all applications concurrently connected to the database. The unit of measure is Pages (4 KB).
- The type is int.

#### Log Buffer Size(Pages)

- Allows you to specify how much of the database heap (defined by the dbheap parameter) that you want to use as a buffer for log records before you write these records to disk. The unit of measure is Pages (4 KB).
- The type is int.

#### Application Control Heap Size(Pages)

- Specifies the maximum size for the application control shared memory. This parameter is used primarily for sharing information between agents working on the same request. The Unit of Measure is Pages. Each Page size is 4 KB.
- The type is int.

#### Lock Timeout(microSec)

- The time taken in microseconds that a request to lock an object timed-out instead of being granted, since the first database connection.
- The type is int.

Number of Primary Log Files

- Specifies the number of primary log files to be pre-allocated. The primary log files establish a fixed amount of storage allocated to the recovery log files.
- The type is int.

Total Time Database Waited for Locks(microSec)

- The total amount of time that the database waited for locks.
- The type is int.

### **Metrics**

Db2ConnInfo Sample Time

- The sample time.
- The type is timestamp.
- The unit is unspecified.

### **Component: SAP SYS Batch Jobs**

Provides information about the configuration, progress, and performance of batch jobs in the mySAP system.

#### **Dimensions**

Definition Time

- The time stamp for the date and time that the batch job was defined.
- The type is timestamp.

Job Class

- A category for the batch job.
- The type is string.

Defined By

- An identifier for the user who defined the batch job. For example, RSMITH specifies the user who defined the batch job.
- The type is string.

jobDetails Sample Interval End

- The time stamp for the stopping time of the data supplied by the Monitoring Agent for mySAP system.
- The type is timestamp.

Start Time

- The date and time the batch job began.
- The type is timestamp.

Job Number

- A numeric identifier for the batch job. This attribute is being deprecated. Refer to the Job ID attribute.
- The type is int.

jobDetails Sample Interval Start

- The time stamp for the beginning time of the data returned by the Monitoring Agent for mySAP system.
- The type is timestamp.

#### Target Instance

- The name of the application instance where this job is configured to run.
- The type is string.

#### jobDetails Last Changed By

- A text string identifier or user ID for the user who last modified the batch job. For example, SBROWN specifies the name of the user who last changed the batch job.
- The type is timestamp.

#### Other Scheduling Value

- A text string identifier for alternative scheduling values. For example, FIRST JOB indicates the job name for an alternate scheduling type of AfterJob.
- The type is string.

#### Job Name

- A text string identifier or name for the batch job. This attribute provides single-byte character support only. For example, COLLECTOR FOR PERFORMANCE specifies the name of the batch job.
- The type is string.

#### Job ID

- String identifier for a batch job. This attribute replaces the Job Number attribute. For example, 1158100A identifies the number of a batch job.
- The type is string.

#### jobDetails Status

- The number of cancelled jobs in the system.
- The type is string. This is a key dimension.

#### Execution Instance

- The name of the instance where this job actually ran.
- The type is string.

#### jobDetails Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### Scheduled Start Time

- The time stamp for the date and time the batch job is scheduled to begin.
- The type is timestamp.

#### Variant

- Name of the variant within a step.
- The type is string.

#### Scheduled Latest Time

- The time stamp for the date and time after which the job must not run.
- The type is timestamp.

#### sapSysJobDetails Last Changed Time

- The time stamp for the date and time the batch job was most recently modified.
- The type is timestamp.

#### jobDetails System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

#### End Time

- The time stamp for the date and time the batch job stopped.
- The type is timestamp.

#### jobDetails Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### Job Name (Unicode)

- A text string identifier or name for the batch job. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string. This is a key dimension.

#### Other Scheduling Type

- A text string identifier for alternative types of scheduling.
- The type is string.

#### jobDetails SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### Target Host

- An identifier or name for the computer designated as the target host.
- The type is string.

#### jobDetails Periodic

- A text string indicator for how often the batch job is scheduled to run. For example, 02 HOURS indicates the job is scheduled to run every two hours.
- The type is string.

#### Execution Host

- The name of the computer serving as the execution host. For example, agoura1 is the name of the computer serving as the execution host.
- The type is string.

#### jobDetails Client

- An identifier for the execution client. For example, 800 indicates the identifier for the client.
- The type is string.

#### jobDetails System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

### Metrics

#### Duration (mins)

- The calculated run time in minutes. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### Number of Steps

- The sum of the number of steps completed for this job. For example, 9 indicates the number of steps completed for this job.
- The type is int.
- The unit is steps.

#### Delayed (seconds)

- A parameter that calculates the delayed time in seconds.
- The type is int.
- The unit is seconds.

### Component: Data Base Summary

Data Base Summary is a system level data set that provides summary information about an Oracle database used in the mySAP system.

#### Dimensions

##### DBSummary System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

##### Database Server Name

- A text string identifier or name for the database server. Use this attribute to specify the name of the database server. For example, ORACLE indicates the name of the database server.
- The type is string.

##### Data Base Summary Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

##### sapDatabaseSummary Object Type

- The category of the database object, such as, table, index, tablespace, or database. For example, Index indicates the type of database object.



- The type is string. This is a key dimension.

#### DatabaseSummary SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### Database Instance Name

- A text string identifier or name for the database instance. For example, CN1 indicates the name of the database instance.
- The type is string.

#### Analysis Time

- The time stamp for the date and time mySAP collected the sample based on a periodic sample schedule.
- The type is timestamp.

#### sapDatabaseSummary Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### Data Base Summary System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

### **Metrics**

#### Total Free (mb)

- The total amount of free space for the database object, in MB.
- The type is double.
- The unit is MB.

#### Total Used (mb) (Superseded)

- The total amount of space used, in MB, for the database object. For example, 5255653 indicates the amount of space used for the database object.
- The type is double.
- The unit is MB.

#### Total Free (kb)

- The total amount of free space, in KB, for the database object.
- The type is int.
- The unit is KB.

#### Total Free (%)

- The total amount of free space for the database object, expressed as a percentage.
- The type is float.
- The unit is percentage.

#### Freespace Problems

- The number of freespace problems. For example, 3 indicates the number of free space problems.
- The type is int.
- The unit is freespace problems.

#### Total Number

- The total number of database objects.
- The type is int.
- The unit is objects.

#### Total Used (%)

- The total amount of space used, expressed as a percentage, for the database object. For example, 51 indicates the percentage amount of space used for the database object.
- The type is float.
- The unit is percentage.

#### Total Size (kb) (Superseded)

- The total amount of space, in KB, for the database object.
- The type is int.
- The unit is KB.

#### Total Used (mb)

- The total amount of space used, in MB, for the database object.
- The type is double.
- The unit is MB.

#### Total Used (kb)

- The total amount of space used, in KB, for the database object.
- The type is int.
- The unit is KB.

#### Total Used (kb) (Superseded)

- The total amount of space used, in KB, for the database object. For example, 5255653 indicates the amount of space used for the database object.
- The type is int.
- The unit is KB.

#### Missing In DDIC

- The number of objects unaccounted for in the Oracle data dictionary. For example, 2 indicates the number of objects unaccounted for in the data dictionary.
- The type is int.
- The unit is objects.

#### Minimum Free (mb) (Superseded)

- The minimum amount of free space, in MB, for the database object. For example, 1928 indicates the amount of free space for a database object.
- The type is double.
- The unit is MB.

#### Total Size (mb) (Superseded)

- The total amount of space, in MB, for the database object. For example, 1045883 indicates the total amount of space for the database object.
- The type is double.
- The unit is MB.

#### sapDatabaseSummary Minimum Free (kb) (Superseded)

- The minimum amount of free space, in KB, for the database object. For example, 1928 indicates the amount of free space for a database object.
- The type is int.
- The unit is KB.

#### Total Free (mb) (Superseded)

- The total amount of free space for the database object, in MB. For example, 5090163 indicates the total MB of free space for the database object.
- The type is double.
- The unit is MB.

#### sapDatabaseSummary Minimum Free (kb)

- The minimum amount of free space, in KB, for the database object.
- The type is int.
- The unit is KB.

#### Total Size (kb)

- The total amount of space, in KB, for the database object.
- The type is int.
- The unit is KB.

#### Total Size (mb)

- The total amount of space, in MB, for the database object.
- The type is double.
- The unit is MB.

#### Total Free (kb) (Superseded)

- The total amount of free space, in KB, for the database object. For example, 5090163 indicates the total amount of free space, in KB, for the database object.
- The type is int.
- The unit is KB.

#### Missing In Database

- The number of objects unaccounted for in the database. Use this attribute to identify the number of objects unaccounted for. For example, 3 indicates the number of objects unaccounted for in the database.
- The type is int.
- The unit is objects.

#### Minimum Free (mb)

- The minimum amount of free space, in MB, for the database object.
- The type is double.
- The unit is MB.

### **Component: Data Base Detail**

Data Base Detail is a system level data set that provides detailed information about an Oracle database used in the mySAP system. Data Base Detail information can be voluminous, so the number of situations written using this data set must be limited to only what is needed and the frequency must be very low. The majority of the database detail information is obtained from the MONI database in SAP. This information is usually updated only once or twice per day. Therefore, there is no benefit to running situations more than once or twice per day.

#### **Dimensions**

Size (kb) (Superseded)

- The defined space, in KB, of the database object. For example, 52553 indicates the defined space of the database object.
- The type is int.

Data Base Detail System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

sapDatabaseDetail Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

DatabaseDetail System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

Object Name

- A text string identifier or name for the database object. For example, REFERENCE indicates the name of the database object.
- The type is string.

Data Base Detail Status

- The status of the database object, such as online, offline, or unknown.
- The type is string.

Size (kb)

- The defined space, in KB, of the database object.
- The type is int.

sapDatabaseDetail Object Type

- The category of the database object, such as, table, index, tablespace, or database. For example, Database indicates the object type.
- The type is string. This is a key dimension.

Space Critical

- Indicates whether space for a database object has reached a critical stage during the last 24 hours. Space critical means that the object fails the next time it needs to extend space, either because of max-extents, tablespace full, or some other reason.
- The type is string.

#### sapDatabaseDetail Analysis Time

- The time stamp for the date and time mySAP collected the sample based on a periodic sample schedule.
- The type is timestamp.

#### DatabaseDetail SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### Data Base Detail Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

### Metrics

#### Max Next Extent (kb) (Superseded)

- The maximum size allowed for the next extent allocated.
- The type is int.
- The unit is KB.

#### Files

- The number of files in tablespace.
- The type is int.
- The unit is files.

#### Tables And Indices Change (per day)

- The number of tables and indices that have changed during the last 24 hours.
- The type is int.
- The unit is tables and indices.

#### sapDatabaseDetail Minimum Free (kb) (Superseded)

- The minimum amount of free space, in KB, in the database object. For example, 3267656 indicates the minimum amount of free space for the database object.
- The type is int.
- The unit is KB.

#### Tables Indices (Superseded)

- The number of tables and indices in tablespace. For example, 523 indicates the number of tables and indices in tablespace.
- The type is int.
- The unit is tables and indices.

#### Tables And Indices Change (per day) (Superseded)

- The number of tables and indices that have changed during the last 24 hours. For example, 23 indicates the number of tables and indices in table space that have changed per day.
- The type is int.
- The unit is tables and indices.

#### Max Next Extent (kb)

- The maximum size allowed for the next extent allocated.
- The type is int.
- The unit is KB.

#### sapDatabaseDetail Minimum Free (kb)

- The minimum amount of free space, in KB, in the database object.
- The type is int.
- The unit is KB.

#### Size Used (%)

- The percentage of space used by the database object. For example, 13 indicates the percentage of space used by the database object.
- The type is float.
- The unit is percentage.

#### Extents Change (per day)

- The number of changes in the reserved blocks of continuous storage per day.
- The type is int.
- The unit is changes.

#### Size Free (kb) (Superseded)

- The amount of space available, in KB, for the database object. Use this attribute to specify the amount of space available for a database object. For example, 5255656 indicates the amount of space available for the database object.
- The type is int.
- The unit is KB.

#### Extents

- The number of reserved blocks of continuous storage.
- The type is int.
- The unit is blocks.

#### Used Change (per day) (Superseded)

- The amount of change, in KB, in the space used by the database object during the last 24 hours. For example, 78533 indicates the amount of space used per day by the database object.
- The type is int.
- The unit is KB.

#### Size Change (per day) (Superseded)

- The amount of change, in KB, in the space used by the database during the last 24 hours. For example, 5893 indicates the amount of change in the space used by the database object.
- The type is int.

- The unit is KB.

#### Size Used (kb)

- The amount of space, in KB, used by the database object.
- The type is int.
- The unit is KB.

#### Maximum Free (kb)

- The maximum amount of free space, in KB, in the database object.
- The type is int.
- The unit is KB.

#### Tables Indices

- The number of tables and indices in tablespace.
- The type is int.
- The unit is tables and indices.

#### Size Used (kb) (Superseded)

- The amount of space, in KB, used by the database object. For example, 45986 indicates the amount of space used by the database object.
- The type is int.
- The unit is KB.

#### Extents Change (per day) (Superseded)

- The number of changes in the reserved blocks of continuous storage per day. For example, 49 indicates the number of changes per day in the reserved blocks.
- The type is int.
- The unit is changes.

#### Extents (Superseded)

- The number of reserved blocks of continuous storage. For example, 43 indicates the number of reserved blocks of continuous storage.
- The type is int.
- The unit is blocks.

#### Size Change (per day)

- The amount of change, in KB, in the space used by the database during the last 24 hours.
- The type is int.
- The unit is KB.

#### Maximum Free (kb) (Superseded)

- The maximum amount of free space, in KB, in the database object. For example, 3267656 indicates the maximum amount of free space for the database object.
- The type is int.
- The unit is KB.

#### Size Free (kb)

- The amount of space available, in KB, for the database object.

- The type is int.
- The unit is KB.

#### Size Free (%)

- The percentage of free space available for the database object. For example, 48 indicates the percentage of free space available for the database object.
- The type is float.
- The unit is percentage.

#### Used Change (per day)

- The amount of change, in KB, in the space used by the database object during the last 24 hours.
- The type is int.
- The unit is KB.

#### Files (Superseded)

- The number of files in tablespace. For example, 236 indicates the number of files in tablespace.
- The type is int.
- The unit is files.

### **Component: Spool Requests**

Spool Requests is a system level data set that provides information about all spool requests in the mySAP system.

#### **Dimensions**

##### Spool Requests Department

- A text string identifier or name for the current department receiving the output of the request. This attribute provides single-byte character support only. For example, PAYROLL indicates the name of the department receiving the output.
- The type is string.

##### sapSpoolRequests Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

##### sapSpoolReq Output Format (Unicode)

- A text string identifier for the current output format. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

##### SpoolReq System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

##### Spool Requests Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.



#### SpoolReq Spool Title

- A text identifier or name for the spool file. For example, LISTISLP01RSMITH indicates the textual identifier of the spool file.
- The type is string.

#### Spool Requests Creator

- The user ID for the originator of the request. For example, RSMITH indicates the user ID for the originator of the request.
- The type is string.

#### sapSpoolRequests Output Device

- A text string identifier or name for the output destination for the spool request. This attribute provides single-byte character support only. For example, LP01 indicates the output destination for the spool request.
- The type is string.

#### Spool Requests Authorization

- An authority object indicating permission to view a spool request. For example, RSMITH in the profile indicates that the user has permission to view a spool request.
- The type is string.

#### Request Closed

- An indicator showing whether the spool file can be appended.
- The type is string.

#### Spool Requests System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

#### Cover Page

- An indicator showing whether a cover page was requested.
- The type is string.

#### sapSpoolReq Output Format

- A text string identifier for the current output format. This attribute provides single-byte character support only. For example, X\_65\_255 indicates the current output format.
- The type is string.

#### sapSpoolRequests Department (Unicode)

- A text string identifier or name for the current department receiving the output of the request. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### SpoolReq Sample Interval End

- The time stamp for the stopping time of the data supplied by the Monitoring Agent for mySAP.
- The type is timestamp.

#### SpoolReq Sample Interval Start

- The time stamp for the beginning time of the data supplied by the Monitoring Agent for mySAP.
- The type is timestamp.

#### sapSpoolRequests Output Device (Unicode)

- A text string identifier or name for the output destination for the spool request. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string. This is a key dimension.

#### Create Time

- The time stamp for the date and time the request was created.
- The type is timestamp.

#### Delete Time

- The time stamp for the date and time after which you can delete the spool file.
- The type is timestamp.

#### Delete After Print

- An indicator showing whether to delete or keep the spool file after printing.
- The type is string.

#### SpoolReq SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### Spool Requests Client

- The identifier or number for the originating client. For example, 800 indicates the client.
- The type is string.

#### Spool Requests Recipient

- A text string identifier or name for the current recipient of the request. For example, RSMITH indicates the name of the current recipient of the request.
- The type is string.

#### SpoolReq Spool Number

- A numeric identifier for the spool file. For example, 31808 indicates the numeric identifier for the spool file.
- The type is int.

### **Metrics**

#### Total Print Requests

- The total number of print requests for this spool request. For example, 3 indicates the total number of print requests for this spool request.
- The type is int.
- The unit is requests.

#### Spool Requests Size

- The size in number of pages available for the spool request. For example, 14638 indicates the number of pages available for the request.
- The type is int.
- The unit is pages.

#### sapSpoolRequests Number of Copies

- The number of copies requested. For example, 3 indicates the number of copies requested.
- The type is int.
- The unit is copies.

#### sapSpoolReq Processed Print Requests

- The total number of processed print requests. For example, 15 indicates the number of processed print requests.
- The type is int.
- The unit is requests.

#### sapSpoolRequests Error Print Requests

- The total number of print requests with errors. For example, 1 indicates the number of print requests with errors.
- The type is int.
- The unit is requests.

### **Component: Transport Requests**

Transport Requests is a system level data set that provides information about all transport requests the mySAP system.

#### **Dimensions**

##### Number (610)

- A text string identifier for the transport request.
- The type is string. This is a key dimension.

##### Parent Number (610)

- A text string identifier for the parent request.
- The type is string.

##### Transport Requests Owner

- A text string identifier or user ID for the owner of the request. For example, RSMITH indicates the user ID for the owner of the request.
- The type is string.

##### Parent Number

- A text string identifier for the parent request. For example, CANKSAV300 indicates an identifier for the parent request.
- The type is string.

##### Transport Requests Number

- A text string identifier for the transport request. For example, CANKSAV300 indicates an identifier for the transport request.
- The type is string.

#### TransportReq Sample Interval Start

- The time stamp for the beginning time of the data supplied by the Monitoring Agent for mySAP.
- The type is timestamp.

#### TransportReq Sample Interval End

- The time stamp for the stopping time of the data supplied by the Monitoring Agent for mySAP.
- The type is timestamp.

#### Transport Requests Type

- The category of the request.
- The type is string.

#### TransportReq System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### Transport Requests Description

- Descriptive text associated with the request. This attribute provides single-byte character support only. For example, Initial Test Transport describes the request.
- The type is string.

#### Transport Requests Status

- The status of the request.
- The type is string.

#### Transport Requests System Name

- The SAP System Identifier (SID) for the mySAP system that you are monitoring. For example, PRD.
- The type is string.

#### Transport Requests Category

- A text string identifier for the Workbench or Customizing category. The Workbench category of requests is associated with changes to planning and business rules. Customizing requests includes modifications to ABAP code or function modules. This attribute provides single-byte character support only. For example, Workbench indicates the workbench category of request.
- The type is string.

#### Highest Return Code

- The highest step return code. Possible values include the following -1 = *blank* 0 = Perfect 4 = Warning 8 = Error 12 = Severe error.
- The type is int.

#### sapTransportRequests Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### Transport Requests Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### Last Changed Time

- The time stamp for the date and time the request was most recently changed.
- The type is timestamp.

#### Import Systems

- A text string identifier for the target systems, or the systems to which the request has been imported. For example, CN1 indicates an identifier for the target system.
- The type is string.

#### Source System

- A text string identifier for the source system, or the system where the request was created. For example, PRD indicates an identifier for the source system.
- The type is string.

#### Description (Unicode)

- Descriptive text associated with the request. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### TransportReq SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### Category (Unicode)

- A text string identifier for the Workbench or Customizing category. The Workbench category of requests is associated with changes to planning and business rules. Customizing requests includes modifications to ABAP code or function modules. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### Source Client

- A text string identifier for the source client where the request was created. For example, 800 indicates an identifier for the source client.
- The type is string.

#### Import Clients

- A text string identifier for the target client to which the request has been imported. For example, 012 indicates an identifier for the target client.
- The type is string.

### **Metrics**

#### Import Count

- The count associated with the import.
- The type is int.

- The unit is import count.

### **SAP Instance**

Information about each mySAP instance. At the instance level, it provides configuration information about one mySAP instance. See the historical data collection section for information about historical data collection for attributes in this data set, including attributes for which data is not collected.

### **Component: File Systems**

File Systems is an instance level data set that provides information about file systems and directory structures used in a mySAP instance.

#### **Dimensions**

FileSystems Sample Time

- The time stamp for the date and time the agent collected data from mySAP system.
- The type is timestamp.

FileSystems File Systems Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

FileSystems Name

- A text string identifier or name for the file system. This attribute provides single-byte character support only. For example, L indicates the name of the file system.
- The type is string.

FileSystems System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

FileSystems Relative Hour

- The sample time.
- The type is string.

FileSystems Message (Unicode)

- Descriptive text indicating the status of the file system. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment. The following values are possible: empty, emptying, rapidly emptying, slowly filling, rapidly filling, slowly, full static.
- The type is string.

FileSystems Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

FileSystems Message

- Descriptive text indicating the status of the file system. This attribute provides single-byte character support only. For example, Static indicates the status of the file system. The following values are possible: empty, emptying rapidly, emptying slowly, filling rapidly, filling slowly, full static.
- The type is string.

#### FileSystems SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### FileSystems Capacity (mb) (Superseded)

- The allocated size, in megabytes, of the file system. For example, 4083 indicates the allocated size of the file system.
- The type is double.

#### FileSystems Name (Unicode)

- A text string identifier or name for the file system. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string. This is a key dimension.

#### FileSystems Instance Name

- The name of the application instance you are monitoring. For example, ddrum2\_PRD\_00 is the name of the application instance you are monitoring.
- The type is string.

#### FileSystems Operating System

- Type of operating system.
- The type is string.

#### FileSystems Capacity (mb)

- The allocated size, in megabytes, of the file system.
- The type is double.

#### FileSystems System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

### **Metrics**

#### FileSystems INodes

- The Total number of Filesystem INodes.
- The type is int.
- The unit is inodes.

#### FileSystems INodes Used (%)

- Percentage of Filesystem INodes used. For UNIX only.
- The type is float.
- The unit is percentage.

#### FileSystems Size Free (mb)

- The amount of space, in megabytes, available in the file system.
- The type is double.
- The unit is MB.

#### FileSystems INodes (Superseded)

- Total number of Filesystem INodes. For UNIX only.
- The type is int.
- The unit is inodes.

#### FileSystems Size Free (mb) (Superseded)

- The amount of space, in megabytes, available in the file system. For example, 108 indicates the amount of space available in the file system.
- The type is double.
- The unit is MB.

#### FileSystems INodes Used (Superseded)

- Filesystem INodes used. For UNIX only.
- The type is int.
- The unit is inodes.

#### FileSystems Size Used (mb)

- The amount of space, in megabytes, used in the file system.
- The type is double.
- The unit is MB.

#### FileSystems Size Used (%)

- The amount of space, expressed as a percentage, used in the file system. For example, 97 indicates the percentage of space used in the file system.
- The type is float.
- The unit is percentage.

#### FileSystems Full Forecast (days)

- The number of days the system estimates that it will take for the file system to become full based on calculations for increased usage during the last 24 hours. A value of -1 indicates that there is no data at this time. This field only displays data when there is an increase in file system usage during the last 24 hours.
- The type is int.
- The unit is days.

#### FileSystems Size Used (mb) (Superseded)

- The amount of space, in megabytes, used in the file system. For example, 3978 indicates the amount of space used in the file system.
- The type is double.
- The unit is MB.

### **Component: Number Range Buffer Summary**

Number Range Buffer Summary is an instance level data set that provides summary and statistical information about the Number Range Buffer used in a mySAP instance.

#### **Dimensions**

##### NumRangeBuffSummary Logon Parameters

- A reserved field for holding execution parameters for KSAR3.



- The type is string.

#### NumRangeBuffSummary Sample Time

- The time stamp for the date and time the agent collected the data from mySAP.
- The type is timestamp.

#### NumRangeBuffSummary Number Range Buffer Summary Description

- The dummy field for the Description column in portrait mode.
- The type is string.

#### NumRangeBuffSummary Number Range Buffer Summary System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

#### NumRangeBuffSummary SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### NumRangeBuffSummary Number Range Buffer Summary System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### NumRangeBuffSummary Number Range Buffer Summary Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### NumRangeBuffSummary Instance Name

- The name of the application instance you are monitoring. For example, ddrum2\_PRD\_00 is the name of the application instance you are monitoring.
- The type is string. This is a key dimension.

#### NumRangeBuffSummary Number Range Buffer Summary Value

- The dummy field for the Value column in portrait mode.
- The type is string.

### Metrics

#### NumRangeBuffSummary Buffer Calls

- The total number of calls to the number range buffer. For example, 78 indicates the number of calls to the number range buffer.
- The type is int.
- The unit is calls.

#### NumRangeBuffSummary Server Responses Less Than 1 msec

- The total number of server responses less than 1 millisecond. For example, 26 indicates the number of buffer responses that are less than 1 millisecond.
- The type is int.
- The unit is responses.

#### NumRangeBuffSummary Server Responses Less Than 50 msec

- The total number of server responses less than 50 milliseconds and greater than 1 millisecond. For example, 54 indicates the number of buffer responses that are less than 50 milliseconds and greater than 1 millisecond.
- The type is int.
- The unit is responses.

#### NumRangeBuffSummary Server Calls

- The number of calls to the server for the number range buffer. For example, 3 indicates the number of calls to the number range server.
- The type is int.
- The unit is calls.

#### NumRangeBuffSummary Buffer Size (kb)

- The allocated buffer size in KB. For example, 669354 indicates the number of KB allocated to the buffer.
- The type is int.
- The unit is KB.

#### NumRangeBuffSummary Buffer Responses 1 msec or Greater

- The total number of buffer responses that are 1 millisecond or greater. For example, 43 indicates the number of buffer responses that are greater than 1 millisecond.
- The type is int.
- The unit is responses.

#### NumRangeBuffSummary Current Entries

- The current number of entries in the number range buffer. For example, 43 indicates the current number of entries in the number range buffer.
- The type is int.
- The unit is entries.

#### NumRangeBuffSummary Number Range Buffer Summary Conflicts

- The total number of number range buffer conflicts. For example, 6 indicates the number of number range buffer conflicts.
- The type is int.
- The unit is conflicts.

#### NumRangeBuffSummary Current Indexes

- The current number of indexes in the number range buffer. For example, 12 indicates the current number of indexes in the number range buffer.
- The type is int.
- The unit is indexes.

#### NumRangeBuffSummary Number Range Buffer Summary Timeouts

- The number of timeouts to the number range buffer. For example, 3 indicates the number of timeouts to the number range buffer.
- The type is int.
- The unit is timeouts.

#### NumRangeBuffSummary Max Entries

- The maximum number of entries in the number range buffer. For example, 1000 indicates the maximum number of entries in the number range buffer.
- The type is int.
- The unit is entries.

#### NumRangeBuffSummary Database Calls

- The number of calls to the database for number ranges. For example, 32 indicates the number of calls to the database for number ranges.
- The type is int.
- The unit is calls.

#### NumRangeBuffSummary Max Indexes

- The maximum number of indexes in the number range buffer. For example, 500 indicates the maximum number of indexes in the number range buffer.
- The type is int.
- The unit is indexes.

#### NumRangeBuffSummary Buffer Responses Less Than 1 msec

- The total number of buffer responses that are less than 1 millisecond and greater than 50 microseconds. For example, 26 indicates the number of buffer responses that are less than 1 millisecond and greater than 50 microseconds.
- The type is int.
- The unit is responses.

#### NumRangeBuffSummary Server Responses 50 msec or Greater

- The number of server responses that are 50 milliseconds or greater. For example, 22 indicates the number of server responses that 50 milliseconds or greater.
- The type is int.
- The unit is responses.

#### NumRangeBuffSummary Buffer Responses Less Than 50 usec

- The total number of buffer responses less than 50 microseconds. For example, 54 indicates the number of buffer responses that are less than 50 microseconds.
- The type is int.
- The unit is responses.

#### NumRangeBuffSummary Get Calls

- The number of get calls to the number range buffer. For example, 78 indicates the number of get calls to the number range buffer.
- The type is int.
- The unit is calls.

### **Component: SAP INS Service Response Time**

Information about the services running in a mySAP instance. These services include batch, dialog, enqueue, gateway, message, spool, and update.

#### **Dimensions**

service Response Time serviceResponseTime Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

service Response Time serviceResponseTime System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

service Response Time serviceResponseTime Sample Interval Start

- The starting time of the data supplied by the SAP agent.
- The type is timestamp.

service Response Time Private Mode Entered

- A text string that indicates whether the private address mode was entered.
- The type is string.

service Response Time serviceResponseTime System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

service Response Time serviceResponseTime Instance Name

- The name of the application instance you are monitoring. For example, ddrum2\_PRD\_00 is the name of the application instance you are monitoring.
- The type is string.

service Response Time serviceResponseTime SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

service Response Time serviceResponseTime Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

service Response Time serviceResponseTime Service Type

- The mySAP service category including batch, dialog, enqueue, gateway, message, spool, and update. For example, Dialog indicates that you are monitoring the mySAP dialog service.
- The type is string. This is a key dimension.

service Response Time serviceResponseTime Sample Interval End

- The time stamp for the stopping time of the data supplied by the SAP agent.
- The type is timestamp.

service Response Time serviceResponseTime Service Type Encoded

- The encoded SAP service type.
- The type is string.

## Metrics

service Response Time Min Response Time (ms)

- The minimum amount of time, in milliseconds, elapsed to process a request for this mySAP service.
- The type is int.
- The unit is milliseconds.

#### service Response Time Service Frequency

- The number of times per minute this service was requested during the sample period. For example, 3 indicates that this service was requested three times per minute during the sampling period.
- The type is double.
- The unit is requests/minute.

#### service Response Time serviceResponseTime Avg CPU Time (ms)

- The average amount of time, in milliseconds, the CPU processed instructions for this transaction.
- The type is int.
- The unit is milliseconds.

#### service Response Time Max CPU Time (ms)

- The maximum amount of time, in milliseconds, the CPU processed instructions for this transaction.
- The type is int.
- The unit is milliseconds.

#### service Response Time Service Frequency (Superseded)

- The number of times per minute this service was requested during the sample period. For example, 3 indicates that this service was requested three times per minute during the sampling period.
- The type is double.
- The unit is requests/minute.

#### service Response Time Avg Wait Time (%)

- The average amount of time, expressed as a percentage, an unprocessed step waited in the queue for a free work process. For example, 10 indicates that the amount of time, expressed as a percentage, an unprocessed step waited in the queue for an available work process averaged ten percent during the sampling period.
- The type is int.
- The unit is percent.

#### service Response Time serviceResponseTime Avg Database Request Time (ms)

- The average amount of time, in milliseconds, the database processed this transaction.
- The type is int.
- The unit is milliseconds.

#### service Response Time Min CPU Time (ms)

- The minimum amount of time, in milliseconds, the CPU processed instructions for this transaction.
- The type is int.
- The unit is milliseconds.

service Response Time Min Database Request Time (ms)

- The minimum amount of time, in milliseconds, elapsed for the database to process this transaction.
- The type is int.
- The unit is milliseconds.

service Response Time Max Database Request Time (ms)

- The maximum amount of time, in milliseconds, elapsed for the database to process this transaction.
- The type is int.
- The unit is milliseconds.

service Response Time Max Response Time (ms)

- The maximum amount of time, in milliseconds, elapsed to process a request for this mySAP service.
- The type is int.
- The unit is milliseconds.

service Response Time Max CPU Time (ms) (Superseded)

- The maximum amount of time, in milliseconds, the CPU processed instructions for this transaction. For example, 180 indicates that the maximum amount of time, in milliseconds, the CPU processed instructions for this transaction was 180 milliseconds during the sampling period.
- The type is int.
- The unit is milliseconds.

service Response Time serviceResponseTime Avg Response Time (ms)

- The average amount of time, in milliseconds, elapsed to process a request for this mySAP service.
- The type is int.
- The unit is milliseconds.

service Response Time Max Wait Time (ms)

- The maximum amount of time, in milliseconds, an unprocessed step waited in the queue for a free work process.
- The type is int.
- The unit is milliseconds.

service Response Time serviceResponseTime Avg Wait Time (ms)

- The average amount of time, in milliseconds, an unprocessed step waited in the queue for a free work process.
- The type is int.
- The unit is milliseconds.

service Response Time Min Wait Time (ms)

- The minimum amount of time, in milliseconds, an unprocessed step waited in the queue for a free work process.
- The type is int.

- The unit is milliseconds.

service Response Time Max Response Time (ms) (Superseded)

- The maximum amount of time, in milliseconds, elapsed to process a request for this mySAP service. For example, 203 indicates that the maximum amount of time elapsed to process a request for this mySAP service was 203 milliseconds during the sampling period.
- The type is int.
- The unit is milliseconds.

service Response Time serviceResponseTime Dialog Steps

- Number of dialog steps.
- The type is int.
- The unit is steps.

### **Component: Gateway Connections**

Gateway Connections is an instance level data set that provides information about the connections between a mySAP instance and external systems.

#### **Dimensions**

GatewayConnections Instance Name

- The name of the application instance that you are monitoring. For example, ddrum2\_PRD\_00 is the name of the application instance you are monitoring.
- The type is string.

GatewayConnections Local Logical Unit Name

- The identifier for the local logical unit. For example, drum2 is an example of a local logical unit name.
- The type is string. This is a key dimension.

GatewayConnections Local APPC Version

- The identifier for the local APPC version. For example, 6 specifies Version 6 of the local APPC.
- The type is int.

GatewayConnections SAP Return Code

- The last SAP return code from structure GWY\_CONNAT, field SAPRC, using function GWY\_READ\_CONNECTION\_ATTRIBUTES. For example, 0 indicates the identifier for the last SAP return code.
- The type is int.

GatewayConnections Symbolic Destination Name

- The symbolic destination name. For example, sapgw00 indicates the symbolic destination name.
- The type is string.

GatewayConnections Remote IP Address

- The identifier for the remote TCP/IP address. For example, 10. 58. 9. 12 is an example of a remote TCP/IP address.
- The type is string.

GatewayConnections Connection or Client Number

- The identifier for the connection number. For example, 6 specifies the connection number.
- The type is int.

#### GatewayConnections Remote Host

- The identifier for the name of the computer serving as the remote host. For example, agoura1 is an example of a remote host name.
- The type is string.

#### GatewayConnections Connection or Client Type

- The type of the mySAP gateway client you are using. For example, LOCAL\_R3 specifies the type of mySAP gateway client.
- The type is string.

#### GatewayConnections Remote Logical Unit Name

- The identifier for the remote logical unit. For example, CAN2 is an example of a remote logical unit name.
- The type is string.

#### GatewayConnections Conversation Identifier

- The identifier for the connection conversation. For example, 862335 specifies the connection conversation number.
- The type is string.

#### GatewayConnections Gateway Connections User ID

- The name of the user making use of the connection. For example, RBROWN is the name of the user connected to the Gateway.
- The type is string.

#### GatewayConnections Request Time

- The time stamp for the time of the last request.
- The type is timestamp.

#### GatewayConnections Local IP Address

- The local TCP/IP address. For example, 195. 0. 2. 3 is an example of a local TCP/IP address.
- The type is string.

#### GatewayConnections Trace Level

- The trace detail level. For example, 0 specifies the trace level.
- The type is int.

#### GatewayConnections Remote APPC Version

- The version number for the remote APPC. For example, 6 specifies Version 6 of the remote APPC.
- The type is int.

#### GatewayConnections Remote Transaction Prog. Name

- The name of the remote transaction program. For example, sapdp00 is an example of a remote transaction program name.
- The type is string.



#### GatewayConnections SNA Return Code

- The identifier for the last SNA return code. For example, 0 indicates the identifier for the last return code. The last SNA return code from structure GWY\_CONNAT, field APPCRC, using function GWY\_READ\_CONNECTION\_ATTRIBUTES.
- The type is int.

#### GatewayConnections Gateway Connections Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### GatewayConnections Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### GatewayConnections Local Transaction Prog. Name

- The name of the local transaction program. For example, ksaagent is an example of a local transaction program name.
- The type is string.

#### GatewayConnections SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### GatewayConnections Gateway Connections System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

#### GatewayConnections Local Host

- The identifier for the name of the computer serving as the local host. For example, CAN2 is an example of a local host name.
- The type is string.

#### GatewayConnections Connection Speed

- The speed of the connection on your mySAP Gateway.
- The type is string.

#### GatewayConnections Gateway Connections Sample Time

- The time stamp for the date and time the agent collected the data from mySAP.
- The type is timestamp.

#### GatewayConnections Local IP Address (v4/v6)

- The local TCP/IP address. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

#### GatewayConnections Gateway Connections Status

- The status of the mySAP Gateway connection. For example, CONNECTED indicates the connection to the gateway is active.
- The type is string.

#### GatewayConnections Remote IP Address (v4/v6)

- The identifier for the remote TCP/IP address. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

#### GatewayConnections In Use

- Indicator of whether or not the connection is in use.
- The type is string.

#### GatewayConnections System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### GatewayConnections Registration Status

- The registration status for the system connection. For example, UNUSED specifies the registration status for the system connection.
- The type is string.

### Metrics

#### GatewayConnections Number of Connections

- The number of connections on your mySAP Gateway. For example, 14 specifies the number of connections.
- The type is int.
- The unit is connections.

### Component: SAP INS Operating System Performance

Information about the operating system on which a mySAP instance is running. The SAP OS collector must be running on the mySAP instance for data to be returned for these attributes.

### Dimensions

#### OSPerf Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### OSPerf SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### OSPerf System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

#### OSPerf Dummy Field for Value

- The dummy field for the Value column in portrait mode.
- The type is string.

#### OSPerf Sample Time

- The time stamp for the date and time the agent collected the data from mySAP.
- The type is timestamp.

#### OSPerf Description Dummy Field

- The dummy field for the Description column in portrait mode.
- The type is string.

#### OSPerf Instance Name

- The name of the application instance you are monitoring. For example, ddrum2\_PRD\_00 is the name of the application instance you are monitoring.
- The type is string. This is a key dimension.

#### OSPerf Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### OSPerf System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

### Metrics

#### OSPerf Physical Memory (kb)

- The total amount of physical memory (RAM). For example, 131136 indicates the total amount of physical memory, in KB.
- The type is int.
- The unit is kilobytes.

#### OSPerf Pages In (sec)

- The number of pages read from disk to update memory references to pages that were not previously referenced during the last second. For example, 0 indicates that no pages were read from disk to update memory references.
- The type is int.
- The unit is pages/second.

#### OSPerf Load Average (1 min)

- The average computing burden the system carried during the last 60 seconds. For example, 0.08 indicates the average computing burden the system carried.
- The type is double.
- The unit is burden.

#### OSPerf Idle CPU Utilization (%)

- The amount of time the CPU is not processing instructions, expressed as a percentage. For example, 93 indicates that the CPU is idle 93 percent of the time it is available.
- The type is int.
- The unit is percent.

#### OSPerf LAN Packets In (sec) (Superseded)

- The number of units, known as packets, that were transferred from the LAN to mySAP during the last second. For example, 2 indicates the number of packets transferred per second.
- The type is int.
- The unit is packets.

#### OSPerf Pages Out (sec)

- The number of modified pages written to disk during the last second. For example, 3 indicates the number of pages written to disk per second.
- The type is int.
- The unit is pages/second.

#### OSPerf Load Average (5 min)

- The average computing burden the system carried during the last five minutes. For example, 0.09 indicates the average computing burden the system carried.
- The type is double.
- The unit is burden.

#### OSPerf User CPU Utilization (%)

- The percentage of CPU used by user tasks.
- The type is int.
- The unit is percent.

#### OSPerf System CPU Utilization (%)

- The percentage of CPU used by system services.
- The type is int.
- The unit is percent.

#### OSPerf Swap Space Free (%)

- The percentage of swap space available. For example, 78 indicates that 78% of swap space is available on this instance.
- The type is int.
- The unit is percent.

#### OSPerf KB Paged Out (sec)

- The number of KB paged out per second.
- The type is int.
- The unit is kilobytes/second.

#### OSPerf LAN Packets In (sec)

- The number of units, known as packets, that were transferred from the LAN to mySAP during the last second.
- The type is int.
- The unit is packets/second.

#### OSPerf Physical Memory Free (kb)

- The amount of physical memory (RAM) available, in KB. For example, 68976 indicates that 67 MB of RAM are available on this instance.
- The type is int.

- The unit is kilobytes.

#### OSPerf LAN Packets Out (sec) (Superseded)

- The number of units, known as packets, that were transferred from mySAP to the LAN during the last second. For example, 2 indicates the number of packets transferred out per second.
- The type is int.
- The unit is packets/second.

#### OSPerf Load Average (15 min)

- The load average during the last 15 minutes.
- The type is double.
- The unit is burden.

#### OSPerf LAN Errors (sec) (Superseded)

- The total number of errors on the LAN during the last second. For example, 2 indicates the total number of errors on the LAN.
- The type is int.
- The unit is errors/second.

#### OSPerf LAN Errors (sec)

- The total number of errors on the LAN during the last second.
- The type is int.
- The unit is errors/second.

#### OSPerf Swap Space (kb)

- The total amount of swap space configured, in KB. For example, 205224 indicates the total amount of swap space configured, in KB.
- The type is int.
- The unit is kilobytes.

#### OSPerf Physical Memory Free (%)

- The percentage of physical memory (RAM) available. For example, 78 indicates that 78% of RAM is available on this instance.
- The type is int.
- The unit is percent.

#### OSPerf Swap Space Free (kb)

- The amount of swap space available, in KB. For example, 411452 indicates that 411 MB of swap space are available on this instance.
- The type is int.
- The unit is kilobytes.

#### OSPerf LAN Collisions (sec) (Superseded)

- The number of times LAN packets could not be delivered because two nodes attempted to send data at the same time. For example, 2 indicates the number of times LAN packets could not be delivered because two nodes attempted to send data at the same time during the last second.
- The type is int.
- The unit is packets.

#### OSPerf LAN Collisions (sec)

- The number of times LAN packets could not be delivered because two nodes attempted to send data at the same time.
- The type is int.
- The unit is occurrences.

#### OSPerf LAN Packets Out (sec)

- The number of units, known as packets, that were transferred from mySAP to the LAN during the last second.
- The type is int.
- The unit is packets./second.

#### OSPerf KB Paged In (sec)

- The number of KB paged in per second.
- The type is int.
- The unit is kilobytes/second.

### **Component: SAP INS Buffer Details**

Provides information about SAP buffers and memory areas in one SAP instance. This data set contains a large number of attributes and it represents buffer utilization. Not all attributes apply to every object reported. When an attribute does not apply to a particular object type, the attribute has a value of -1.

#### **Dimensions**

##### sap Buffers sapBuffers Encoded Name

- Encoded version of the Name attribute.
- The type is string.

##### sap Buffers sapBuffers SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

##### sap Buffers sapBuffers Managed System

- The identifier for this SAP resource.
- The type is string.

##### sap Buffers sapBuffers Instance Name

- Name of the application instance that you are monitoring. For example, ddrum2\_PRD\_00 is the name of the application instance that you are monitoring.
- The type is string.

##### sap Buffers sapBuffers System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the data base server associated with the target SAP system.
- The type is string.

##### sap Buffers sapBuffers Name

- A text string identifier or name for the buffer or memory area. For example, Heap memory indicates the name of a memory area and IRBD Initial Records indicates the name of the buffer.
- The type is string. This is a key dimension.

sap Buffers sapBuffers Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

sap Buffers sapBuffers System Name

- The SAP System Identifier (SID) for the SAP system that you are monitoring. For example, PRD.
- The type is string.

sap Buffers sapBuffers Last Reset

- Time stamp for the most recent date and time that the buffer was cleared out.
- The type is timestamp.

## Metrics

sap Buffers sapBuffers hits

- Amount of times that the requested data was available in the buffer. For example, 17268 indicates the number of times the data was available in the buffer.
- The type is int.
- The unit is hits.

sap Buffers sapBuffersSize Allocated (kb)

- Number of space in KB allocated to the buffer.
- The type is int.
- The unit is KB.

sap Buffers sapBuffers Size Free (kb)

- Number of buffer space or memory area available, in KB. For example, 4836 indicates the amount of buffer space available.
- The type is int.
- The unit is KB.

sap Buffers sapBuffers Size Used (kb)

- Number of buffer space used, in KB. Use this attribute to specify the amount of buffer space used. For example, 629 indicates the amount of buffer space used.
- The type is int.
- The unit is KB.

sap Buffers sapBuffers Inserts

- Number of buffer inserts. For example, 28 indicates the number of buffer inserts.
- The type is int.
- The unit is inserts.

sap Buffers sapBuffers Size Reserved (kb)

- Size reserved by SAP for internal buffer management. The value is Size Allocated minus Size Used and Size Free.
- The type is int.
- The unit is KB.

#### sap Buffers sapBuffers DB Accesses

- Number of times the database was accessed when the requested data was not available in the buffer. For example, 254 indicates the number of times the database was accessed.
- The type is int.
- The unit is database access.

#### sap Buffers sapBuffers Hitratio (%)

- An identifier, expressed as a percentage, indicating the percentage of requests that were satisfied from the buffer. The percentage is calculated as follows:  $(\text{buffer\_hits} * 100) / \text{buffer\_requests}$  and it must be close to 100%. For example, 99.37 indicates the percentage of requests that were satisfied from the buffer.
- The type is double.
- The unit is percent.

#### sap Buffers sapBuffers Deletes

- Number of buffer deletes. For example, 9 indicates the amount of buffer deletes.
- The type is int.
- The unit is buffer deletes.

#### sap Buffers sapBuffers Size In Memory (kb)

- A memory size metric specific to roll area, page area, and extended memory. A disk-size metric specific to roll area and page area -1 = N/A.
- The type is int.
- The unit is KB.

#### sap Buffers sapBuffersSize On Disk (kb)

- A disk-size metric specific to roll area and page area. The valid format is a 8-byte integer.
- The type is int.
- The unit is KB.

#### sap Buffers sapBuffers Changes

- Number of buffer updates. For example, 9 indicates the number of buffer updates.
- The type is int.
- The unit is buffer updates.

#### sap Buffers sapBuffers Max Used (kb)

- Metric specific to roll area, page area, extended memory and heap.
- The type is int.
- The unit is KB.

#### sap Buffers sapBuffers Size Free (%)

- Free percentage for buffers and memory areas such as roll, page, extended memory.
- The type is int.
- The unit is percent.

#### sap Buffers sapBuffers Objects In Buffer

- Number of objects in the buffer For example, 189 indicates the number of objects in the buffer.



- The type is int.
- The unit is objects.

#### sap Buffers sapBuffers Directory Allocated

- Maximum amount of objects that the buffer holds because one directory entry is required for each object that the buffer contains. For example, 12289 indicates the amount of directory entries defined for a buffer.
- The type is int.
- The unit is directory entries.

#### sap Buffers sapBuffers Total Resets

- Total number of times that the buffer space was cleared out. Resets occur automatically during system initialization, as well as manually. For example, 9 indicates the number of times that the buffer space was cleared out.
- The type is int.
- The unit is resets.

#### sap Buffers sapBuffers Directory Used (%)

- Percentage of the directory that was used.
- The type is int.
- The unit is percent.

#### sap Buffers sapBuffers Requests

- Number of buffer requests.
- The type is int.
- The unit is requests.

#### sap Buffers sapBuffers Directory Free (%)

- Percentage of the directory that is free.
- The type is int.
- The unit is percent.

#### sap Buffers sapBuffers Objects Swapped

- Number of objects swapped in the buffer For example, 3 indicates the number of objects swapped in the buffer.
- The type is int.
- The unit is swapped objects.

#### sap Buffers sapBuffers Directory Free

- Number of directory entries that are currently not in use. If the buffer size is large enough, this amount determines the new objects that you can add to this buffer. For example, 12140 indicates the number of directory entries not in use.
- The type is int.
- The unit is free directory entries.

#### sap Buffers sapBuffers Misses

- Number of times that the requested data was not available in the buffer. For example, 468 indicates the number of times the requested data was not available in the buffer.

- The type is int.
- The unit is misses.

#### sap Buffers sapBuffers Size Reserved (%)

- Percentage reserved by SAP for internal buffer management. The value is Size Reserved divided by Size Allocated.
- The type is int.
- The unit is percent.

#### sap Buffers sapBuffers Frames Swapped

- Number of frames swapped in the buffer. For example, 1 indicates the number of frames swapped in the buffer.
- The type is int.
- The unit is frames swapped.

#### sap Buffers sapBuffers DB Accesses Saved

- Number of times the database accesses were saved. Database accesses occur when the requested data is not available in the buffer. For example, 57456 indicates the number of times that the database accesses were saved.
- The type is int.
- The unit is saved database accesses.

#### sap Buffers sapBuffers DB Access Quality (%)

- An indicator expressed as a percentage to indicate the percentage of requests that were satisfied from the buffer. This percentage must be close to 100%, and is calculated as follows:  $(db\_accesses\_saved * 100) / (db\_accesses + db\_accesses\_saved)$ . For example, 99.37 indicates the percentage of requests that were satisfied.
- The type is double.
- The unit is percent.

#### sap Buffers sapBuffers Max Used (%)

- Metric specific to roll area, page area, and extended memory.
- The type is int.
- The unit is percent.

#### sap Buffers sapBuffers Directory Used

- Number of directory entries that are currently in use, which is the number of objects currently in the buffer. For example, 149 indicates the amount of directory entries currently in use.
- The type is int.
- The unit is used directory entries.

#### sap Buffers sapBuffers Size Used (%)

- Percentage used for buffers and memory areas such as roll and page.
- The type is int.
- The unit is percent.

### **Component: SAP Instance Details**

Provides information about SAP Instance.

## Dimensions

### PI Central Instance Name

- The name of the central instance application server that is configured for this mySAP system.
- The type is string.

### PI Operation Mode

- A text string identifier or name for the current operation mode of the system. For example, Private indicates the current operation mode of the system. This attribute provides single-byte character support only.
- The type is string.

### PI Database Release

- Release associated with the database.
- The type is string.

### PI Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

### PI Dialog Service Configured

- A Yes/No switch to indicate if the dialog service is configured.
- The type is string.

### PI Instance Stop Time

- The time stamp for the date and time the application instance stopped.
- The type is timestamp. This is a key dimension.

### PI\_Instance Host Name

- The name of the physical system, without the domain, on which this application server resides. For example, Insthost is the name of the application instance you are monitoring.
- The type is string. This is a key dimension.

### PI System Release

- The release number for the level of software installed on this application server. For example, 640 indicates the level of software installed in the SAP mySAP system you are monitoring.
- The type is string. This is a key dimension.

### PI System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

### PI Value

- The dummy field for the Value column in portrait mode.
- The type is string.

### PI Message Service Configured

- A Yes/No switch to indicate if the message server is configured.
- The type is string.

#### PI SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### PI System Description

- A user-provided description of this application instance as defined in the mySAP system transport table. This attribute provides single-byte character support only.
- The type is string. This is a key dimension.

#### PI Assigned Update Instance

- The name of the application server assigned to a specific update server. For example, Updinst\_SY1\_00 is the instance configured with the mySAP update service for this application instance.
- The type is string.

#### PI Database Host IP Address

- The IP address of the physical system on which the database instance resides. This value is the same for all instances of a mySAP system. For example, 170. 106. 1. 1 is the IP address for the database host in the mySAP system you are monitoring.
- The type is string. This is a key dimension.

#### PI Instance Host IP Address

- The IP address of the physical system on which the application instance resides. For example, 170. 106. 1. 11 is the IP address of the physical system on which the application instance you are monitoring resides.
- The type is string. This is a key dimension.

#### PI Gateway Service Configured

- A Yes/No switch to indicate if the gateway service is configured.
- The type is string.

#### PI Central Instance

- A Yes/No switch to indicate if the application server is the central instance. This attribute can be useful when tailoring a situation.
- The type is string.

#### PI Database Host IP Address (v4/v6)

- The IP address of the physical system on which the database instance resides. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

#### PI System Number

- The number assigned to this application server instance. For example, 01 is the number of the mySAP instance you are monitoring.
- The type is string. This is a key dimension.

#### PI Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string. This is a key dimension.

#### PI Enqueue Service Configured

- A Yes/No switch to indicate if the enqueue service is configured.
- The type is string.

#### PI Instance Start Time

- The time stamp for the date and time the application instance started.
- The type is timestamp. This is a key dimension.

#### PI Instance Name

- The name of the application server.
- The type is string. This is a key dimension.

#### PI Description

- The dummy field for the Description column in portrait mode.
- The type is string.

#### PI Instance Host IP Address (v4/v6)

- The IP address of the physical system on which the application instance resides. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

#### PI System Description (Unicode)

- A user-provided description of this application server instance as defined in the mySAP system transport table. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### PI Database Name

- The name of the database instance defined for this mySAP system. This name is frequently the same as the mySAP SID, and is the same for each instance of a mySAP system. For example, DB4 is the name of the physical system on which the database server resides in the mySAP system you are monitoring.
- The type is string.

#### PI System Start Time

- The time stamp for the date and time the system started.
- The type is timestamp.

#### PI Instance Op Mode State

- The state in which the instance is included in the current operation mode of this application server.
- The type is string.

#### PI Database Host Name

- The name of the host computer running the database instance of a system. For example, DBhost is the name of the database host in the mySAP system you are monitoring.
- The type is string. This is a key dimension.

#### PI Sample Time

- The time stamp for the date and time the agent collected the data from mySAP.
- The type is timestamp.

#### PI Operation Mode (Unicode)

- A text string identifier or name for the current operation mode of the system. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### PI System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### PI Batch Service Configured

- A Yes/No switch to indicate if the batch service is configured.
- The type is string.

#### PI Spool Service Configured

- A Yes/No switch to indicate if the spool service is configured.
- The type is string.

#### PI Configuration String

- The services mask, or string, for this application server. For example, DVEBMGS indicates that the following mySAP services are configured for this instance: D = Dialog V = Update (stands for Verbucher in German) E = Enqueue B = Background M = Message server G = SNA gateway S = Spool.
- The type is string.

#### PI Update Service Configured

- A Yes/No switch to indicate if the update service is configured.
- The type is string.

#### PI Instance Status

- The status of this application instance, either running or not running.
- The type is string. This is a key dimension.

#### PI Update2 Service Configured

- A Yes/No switch to indicate if the Update2 service is configured.
- The type is string.

#### PI Database Type

- Type of database.
- The type is string. This is a key dimension.

### **Metrics**

#### PI Registered Users

- Number of total registered users currently for this SAP system.

- The type is int.
- The unit is users.

#### PI Update2 Queue Percent

- Percentage of the dispatcher queue allotted for an Update2 that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI System Up Duration

- The amount of time, in minutes, that the system has been up. For example, 12 indicates that the system has been up for 12 minutes. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### PI Update2 Running Percent

- Percent of Update2 work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Total Active Users

- Number of active users currently for this server. It includes RFC users and interactive users.
- The type is int.
- The unit is users.

#### PI Update Complete Percent

- Percent of Update work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Dialog Running Percent

- Percent of Dialog work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Batch Complete Percent

- Percent of Batch work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI\_Instance Down Duration

- The amount of time, in minutes, an application instance has been down. For example, 12 indicates that a particular instance has been down for 12 minutes. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### PI Total RFC Sessions

- The total number of RFC sessions.
- The type is int.
- The unit is RFC sessions.

#### PI Dialog Waiting Percent

- Percent of Dialog work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### PI Batch Stopped Percent

- Percent of Batch work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Enqueue Waiting Percent

- Percent of Enqueue work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### PI Spool Complete Percent

- Percent of Spool work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI RFC Users

- Number of RFC users currently for this server.
- The type is int.
- The unit is users.

#### PI Dialog Processes

- The number of dialog processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Spool Processes

- The number of spool processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Update Stopped Percent

- Percent of Update work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Update2 Stopped Percent

- Percent of Update2 work processes in the Stopped state.
- The type is int.



- The unit is percent.

#### PI Active Users

- The current number of users logged on to this application instance. For example, 47 indicates the number of users currently logged on to the instance you are monitoring.
- The type is int.
- The unit is users.

#### PI Enqueue Queue

- The number of tasks in the dispatch queue waiting for an Enqueue work process.
- The type is int.
- The unit is tasks.

#### Instances ConnectionFailed

- The total number of instances that have lost connection in the system.
- The type is int.
- The unit is instances.

#### PI Spool Queue Percent

- Percentage of the dispatcher queue allotted for Spool that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Update Processes

- The number of update processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Spool Queue

- The number of tasks in the dispatch queue waiting for a Spool work process.
- The type is int.
- The unit is tasks.

#### PI Spool Stopped Percent

- Percent of Spool work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Spool Running Percent

- Percent of Spool work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Batch Running Percent

- Percent of Batch work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Dialog Queue Percent

- Percentage of the dispatcher queue allotted for Dialog that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Instances Running

- The total number of instances that are running in this system. For example, 15 indicates that 15 instances you are monitoring are running.
- The type is int.
- The unit is instances.

#### PI Spool Waiting Percent

- Percent of Spool work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### Instances Passive

- The total number of instances that are in passive state in the system.
- The type is int.
- The unit is instances.

#### PI Update Running Percent

- Percent of Update work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Update Queue

- The number of tasks in the dispatch queue waiting for an Update work process.
- The type is int.
- The unit is tasks.

#### PI Dialog Stopped Percent

- Percent of Dialog work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Instances Down

- The total number of application instances that are down in this system. This values are only reported for instances defined in an operation mode profile. For example, 3 indicates that 3 of the instances you are monitoring are not running.
- The type is int.
- The unit is instances.

#### PI Update2 Complete Percent

- Percent of Update2 work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Update2 Waiting Percent

- Percent of Update2 work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### PI Interactive Users

- Number of interactive (GUI) users currently for this server.
- The type is int.
- The unit is users.

#### PI NoWP Queue

- The number of tasks in the dispatch queue waiting to be processed by the dispatcher itself or some other system service.
- The type is int.
- The unit is tasks.

#### PI Batch Job Queue

- The number of batch jobs in Ready state.
- The type is int.
- The unit is jobs.

#### PI\_Instance Up Duration

- The amount of time, in minutes, an application instance has been up in this system. For example, 12 indicates that a particular instance has been up for 12 minutes. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### PI Update2 Processes

- Number of Update2 work processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Enqueue Complete Percent

- Percent of Enqueue work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Database Port

- Database Port.
- The type is int.
- The unit is port.

#### PI Dialog Queue

- The number of tasks in the dispatch queue waiting for a Dialog work process.
- The type is int.
- The unit is tasks.

#### PI Enqueue Queue Percent

- Percentage of the dispatcher queue allotted for Enqueue that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Dialog Complete Percent

- Percent of Dialog work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Batch Processes

- The number of batch processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Update Queue Percent

- Percentage of the dispatcher queue allotted for Update that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Enqueue Processes

- Number of enqueue work processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Total External Sessions

- The total number of user sessions (GUI and RFC).
- The type is int.
- The unit is user sessions.

#### PI Batch Waiting Percent

- Percent of Batch work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### PI Update2 Queue

- The number of tasks in the dispatch queue waiting for an Update2 work process.
- The type is int.
- The unit is tasks.

#### PI Enqueue Stopped Percent

- Percent of Enqueue work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Enqueue Running Percent

- Percent of Enqueue work processes in the Running state.

- The type is int.
- The unit is percent.

#### PI Total GUI Sessions

- The total number of non-APPC-TM GUI sessions.
- The type is int.
- The unit is GUI sessions.

#### PI Update Waiting Percent

- Percent of Update work processes in the Waiting state.
- The type is int.
- The unit is percent.

### **Component: SAP INS Alerts**

Information about CCMS and mySAP Agent alerts occurring in a mySAP instance. CCMS alerts are similar to IBM Tivoli Monitoring situations in that they alert you to conditions in which a monitored valued has exceeded a threshold value.

#### **Dimensions**

##### Alert Extended Alert Unique Identifier

- The extended alert unique identifier that is used to close an alert in a SAP system.
- The type is string.

##### Alert Object Name

- The MTE object name, a text string. This attribute applies to CCMS alerts only.
- The type is string.

##### Alert sapAlerts SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

##### Alert Identifier

- A unique identifier assigned by the SAP agent that represents the alert type and subtype. Use this numeric value or range of values to identify or exclude an alert. For example, 517.
- The type is int. This is a key dimension.

##### Alert Default Period

- The new default period will be calculated from where clause generatednby this call.
- The type is int.

##### Alert Severity

- Actual alert severity value from the SAP system.
- The type is int.

##### Alert Reset Alert Inside R/3 Action

- Reset alert inside R/3 action.
- The type is string.

##### Alert sapAlerts System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

#### Alert MTE Class

- A text string for the monitoring tree element in CCMS with which this alert is associated.
- The type is string.

#### Alert Occurrence Time GMT

- The time at which the alert occurred in Greenwich mean time.
- The type is timestamp.

#### Alert sapAlerts Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### Alert Occurrence Time

- The time stamp for the date and time that an alert or range of alerts occurred.
- The type is timestamp. This is a key dimension.

#### Alert sapAlerts Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### Alert Unique Identifier

- The alert unique identifier that is used to close an alert in a SAP system.
- The type is int.

#### Alert Text

- The text associated with an alert generated by mySAP system. This attribute provides single-byte character support only. For example, NO BACKUPS ON RECORD indicates that no backup was detected. For CCMS alerts, this attribute contains a concatenation of all of the texts from the branches of the CCMS alert tree. This is the whole alert tree for the single alert in one attribute.
- The type is string.

#### Alert Msg

- An alert message from the CCMS that provides more details on the reason for the alert.
- The type is string.

#### Alert Closure Time

- The time at which the alert was closed.
- The type is timestamp.

#### Alert Category

- A category associated with an alert, as defined by mySAP system. For example, DATABASE indicates that this alert involves database performance.
- The type is string.

#### Alert Field Name

- The MTE attribute name, a text string. This attribute applies to CCMS alerts only.
- The type is string.

#### Alert Value

- The severity value from the CCMS.
- The type is int.

#### Alert Message (Unicode)

- The text associated with an alert generated by mySAP system. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment. For CCMS alerts, this attribute contains a concatenation of all of the texts from the branches of the CCMS alert tree. This is the whole alert tree for the single alert in one attribute.
- The type is string.

#### Alert sapAlerts Instance Name

- The name of the application instance you are monitoring, a text string. For example, DDRUM2\_PRD\_00.
- The type is string.

#### Alert Status

- The alert status, a number that indicates Open or Acknowledged. This attribute applies to CCMS alerts only.
- The type is int.

#### Alert TID Internal Handle

- Internal handle for TID that is used for the link from the current state view to the alert view.
- The type is string.

#### Alert Raised By

- The system that raised the alert, which is either the SAP agent or mySAP CCMS, a text string value. The following values are included S = SAP C = IBM Tivoli Monitoring.
- The type is string.

#### Alert Monitoring Segment Name

- Name of the monitoring segment.
- The type is string.

#### Alert sapAlerts Sample Time

- The time stamp for the date and time the agent collected the data.
- The type is timestamp.

#### Alert Monitor

- The CCMS Monitor to which this alert belongs, a text string. This attribute applies to CCMS alerts only.
- The type is string.

#### Alert Logon Parameters 1

- Dummy field for holding execution parameters for ksar3.

- The type is string.

#### Alert Logon Parameters 2

- Dummy field for holding execution parameters for ksar3.
- The type is string.

#### Alert sapAlerts System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the database server associated with the target mySAP system.
- The type is string.

#### Alert Monitor Set

- The CCMS Monitor Set to which this alert belongs, a text string. This attribute applies to CCMS alerts only.
- The type is string.

#### Alert Action (610)

- Reset alert inside R/3 action. In ITM 6.1 increase length by 10 to allow for ALUNIQNUM. It now holds keyword RESET plus the MTUID in character format and ALUNIQNUM in character format.
- The type is string.

#### Alert Index

- Internal handle for the alert id.
- The type is string.

### Metrics

#### Alert Critical Alerts

- The number of critical alerts.
- The type is int.
- The unit is alerts.

### Component: SAP INS Work Processes

Information about all work processes running within a mySAP instance.

#### Dimensions

##### work Processes Current Activity

- The current activity of the mySAP work process.
- The type is string.

##### work Processes workProcesses Instance Name

- The name of the application instance you are monitoring. For example, ddrum2\_PRD\_00 is the name of the application instance you are monitoring.
- The type is string. This is a key dimension.

##### work Processes workProcesses System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.



#### work Processes workProcesses SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### work Processes workProcesses Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### work Processes Restart After Error

- A Yes/No switch that indicates whether a process must be restarted automatically after an abnormal termination during its execution.
- The type is string.

#### work Processes Wait Information

- Information supplied by the mySAP system that explains why a process had to wait before executing. This attribute provides single-byte character support only. For example, CMRCV/6066760 is an example of wait information.
- The type is string.

#### work Processes Program (Unicode)

- A text string identifier or name for the program that is currently executing in a work process. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### work Processes workProcesses Sample Time

- The time stamp for the date and time the agent collected the data.
- The type is timestamp.

#### work Processes Filter Field

- Filler field to fix mysterious alignment problem.
- The type is string.

#### work Processes Work Process State

- The current state of the work process.
- The type is string.

#### work Processes workProcesses System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

#### work Processes Identifier

- The identifier for the mySAP process. For example, 6 is the identifier for a particular mySAP work process.
- The type is int.

#### work Processes workProcesses Type

- The type of work process, such as dialog or batch. For example, UPD specifies a work process that executes dialog steps.

- The type is string. This is a key dimension.

#### work Processes Table Name

- The name of the table currently being used by the work process. For example, TADIR is the name of the table currently being used.
- The type is string.

#### work Processes Wait Start Time

- The time stamp for the date and time the process started waiting to execute.
- The type is timestamp.

#### work Processes OS Process Id

- The identifier for the operating system process. For example, 5032 is the number of the operating system process.
- The type is int.

#### work Processes Client

- A text string identifier or name for the client in which the session is running. For example, 800 identifies the name of the client for this session.
- The type is string. This is a key dimension.

#### work Processes Transaction Code

- The identifier for the transaction code. This attribute provides single-byte character support only. For example, FB01 is a transaction code.
- The type is string.

#### work Processes Program Identifier

- A text string identifier or name for the program that is currently executing in a work process. This attribute provides single-byte character support only. For example, SAPETHFB identifies the name of the program associated with this process.
- The type is string.

#### work Processes Wait Information (Unicode)

- Information supplied by the mySAP system that explains why a process had to wait before executing. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### work Processes Transaction Code (Unicode)

- The identifier for the transaction code. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### work Processes Status Reason

- The reason the process stopped.
- The type is string.

#### work Processes Person Name

- The name of the person whose request is being processed. For example, LBROWN is the name of the person using this work process.
- The type is string.

work Processes workProcesses Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

## Metrics

work Processes Database Changes (Superseded)

- The number of database changes, such as deletes, inserts, or updates, that occurred during the execution of a mySAP process. For example, 126 indicates that 126 database changes occurred during the execution of a particular process.
- The type is int.
- The unit is changes.

work Processes Process Roll Size(kb)

- The roll size, in KB, consumed by the process. For example, 114688 is the roll size, in KB, consumed by the mySAP work process.
- The type is int.
- The unit is kilobytes.

work Processes Roll In-Out Count

- The number of roll in actions (where memory is retrieved from roll space), and roll out actions (where memory is temporarily saved to roll space) associated with this current user ID.
- The type is int.
- The unit is occurrences.

work Processes Database Reads Time (ms) (Superseded)

- The amount of time it took, in milliseconds, to perform database reads during the execution of a mySAP process. For example, 1087655 indicates that it took 1,087,655 milliseconds to perform database reads during the execution of a particular mySAP process.
- The type is int.
- The unit is milliseconds.

work Processes Errors

- The number of errors that occurred during the execution of a mySAP process. For example, 03 indicates that 3 errors occurred during the execution of a particular mySAP process.
- The type is int.
- The unit is errors.

work Processes Database Changes

- The number of database changes, such as deletes, inserts, or updates, that occurred during the execution of a mySAP process.
- The type is int.
- The unit is changes.

work Processes Process Private Memory (kb)

- The private memory, in KB, allocated to the process.
- The type is int.
- The unit is kilobytes.

#### work Processes Elapsed Time (secs)

- The amount of time, in seconds, that elapsed during the execution of the current request. For example, 59 indicates that 59 seconds elapsed during the execution of the current request.
- The type is int.
- The unit is seconds.

#### work Processes Roll In-Out Time (ms)

- The amount of time, in milliseconds, spent processing roll ins and roll outs for this mySAP process. For example, 261636 is the amount of time in milliseconds it took to process roll ins and roll outs for this mySAP process.
- The type is int.
- The unit is milliseconds.

#### work Processes Process Page Size(kb)

- The page size, in KB, consumed by the process. For example, 3 is the page size, in KB, consumed by the mySAP process.
- The type is int.
- The unit is kilobytes.

#### work Processes CPU Time (secs)

- The amount of time, in seconds, the CPU spent processing instructions for this mySAP process.
- The type is int.
- The unit is seconds.

#### work Processes Process Total Memory (kb)

- The total amount of private memory, in KB, allocated to the process.
- The type is int.
- The unit is kilobytes.

#### work Processes Database Reads Time (ms)

- The amount of time it took, in milliseconds, to perform database reads during the execution of a mySAP process.
- The type is int.
- The unit is milliseconds.

#### work Processes Database Reads

- The number of database reads that occurred during the execution of an mySAP process. For example, 479 indicates that 479 database reads occurred during the execution of a particular process.
- The type is int.
- The unit is reads.

#### work Processes Database Changes Time (ms)

- The amount of time it took, in milliseconds, to process database changes, such as deletes, inserts, or updates, during the execution of a mySAP process. For example, 374103 indicates that it took 374,103 milliseconds to process certain database changes during the execution of a particular mySAP process.
- The type is int.
- The unit is milliseconds.

### **Component: Active Users**

Active Users is an instance level data set that provides information about users that are currently logged on to a mySAP instance.

#### **Dimensions**

##### ActiveUsers Transaction Code

- The transaction code in which the most recent activity for a user took place. The code identifies each program that can be started from a menu in the mySAP system using a text string. For example, ST03 is the identifier for the mySAP transaction code.
- The type is string.

##### ActiveUsers Session Time

- The time stamp for the date and time of the last session.
- The type is timestamp.

##### ActiveUsers Active Users Time

- The time stamp for the time of the last user activity.
- The type is timestamp.

##### ActiveUsers Session Title (Unicode)

- The screen title of the session, a text string. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

##### ActiveUsers Sample Time

- The time stamp for the date and time the agent collected the data.
- The type is timestamp.

##### ActiveUsers Active Users Client

- A text string identifier or name for the source client session. For example, 800 identifies the name of the client for this session.
- The type is string.

##### ActiveUsers Active Users terminal

- The host name of the terminal running the SAPGUI presentation. Use this text string attribute to specify or exclude a specific terminal. For example, LBROWN is the name of the terminal being used.
- The type is string.

##### ActiveUsers User Key

- The numeric identifier for the memory protection key for the user. For example, 216 is the name of the memory protection key for the user.

- The type is int.

#### ActiveUsers SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### ActiveUsers System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### ActiveUsers IP Address (v4/v6)

- The IP address of the workstation running the SAPGUI presentation. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

#### ActiveUsers Type

- Type of connection for each user who is logged on to the SAP instance. The following values provide information about the connection: 4 = GUI 32 = RFC 202 = PLUGIN 2 = SYSTEM.
- The type is string.

#### ActiveUsers IP Address

- The IP address of the workstation running the SAPGUI presentation. For example, 170. 106. 1. 1 is the IP address.
- The type is string.

#### ActiveUsers Echoed To Session

- The user ID for a different user on this mySAP system. Use this text string attribute to identify sessions being echoed to the SAPGUI screens of other users for the purposes of monitoring, troubleshooting, or training personnel. For example, LBROWN identifies the name of the session echoed.
- The type is string.

#### ActiveUsers Session Title

- The screen title of the session, a text string. For example, ABAP/4 Function Modules is the screen title of the session. This attribute provides single-byte character support only.
- The type is string.

#### ActiveUsers Instance Name

- The name of the application instance you are monitoring. The valid format is a text string. For example, ddrum2\_PRD\_00 is the name of the application instance you are monitoring.
- The type is string.

#### ActiveUsers Active Users User ID

- The name of the user logged on to this session, a text string. For example, LBROWN is the name of the person using this session.
- The type is string. This is a key dimension.

#### ActiveUsers Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.

- The type is string.

ActiveUsers Active Users System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

ActiveUsers Session Number

- The identifier for the user session, a numeric value. For example, 2 is the number of the session.
- The type is int.

ActiveUsers Active Users Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

## Metrics

ActiveUsers User Roll Size (kb)

- The roll size (where user memory is temporarily saved and retrieved from roll space), in KB, allocated to the user, a numeric value. For example, 11468 is the roll size allocated to the user.
- The type is int.
- The unit is KB.

ActiveUsers User Private Memory (kb)

- The private memory, in KB, allocated to the user, a numeric value. For example, 34267 is the private memory allocated to the user.
- The type is int.
- The unit is KB.

ActiveUsers External Sessions

- An integer value for the total number of external (true) sessions. For example, 2 specifies the total number of external sessions.
- The type is int.
- The unit is sessions.

ActiveUsers User Total Memory (kb)

- The total memory, in KB, consumed by the user, a numeric value. For example, 739313 is the total memory consumed by the user.
- The type is int.
- The unit is KB.

ActiveUsers Internal Sessions

- An integer value for the total number of automatically opened internal sessions. For example, 3 specifies the total number of internal sessions.
- The type is int.
- The unit is sessions.

ActiveUsers User Page Size (kb)

- The page size, in KB, consumed by the user, a numeric value. For example, 16384 is the page size consumed by the user.

- The type is int.
- The unit is KB.

### **Component: SAP INS Transaction Performance**

Information about transaction response time and performance characteristics within a mySAP instance.

#### **Dimensions**

transPerf User Name

- The name of the user performing the transaction. For example, RBROWN is the name of the user performing the transaction.
- The type is string. This is a key dimension.

transPerf Instance Name

- The name of the application instance you are monitoring. For example, ddrum2\_PRD\_00 is the name of the application instance you are monitoring.
- The type is string.

transPerf Description (Unicode)

- The program name, transaction code, business application, or user ID description. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment. Description is a language dependent description for the transaction code or program or how the unit of work was started. See the examples above for the Description attribute.
- The type is string.

transPerf Business Application Name

- The name of the business application name or of the sub-application name. This attribute provides single-byte character support only. For example, FI01 is the name of the business application you are monitoring.
- The type is string.

transPerf Program or Tran Code (Unicode)

- The unit of work that you started. It is determined by SAP code. For example
  - Running the FB01 transaction creates a value of FB01 for Transaction Code or Program.
  - Running the RSPFPAR program through transaction SE38 creates a value of RSPFPAR for Transaction code or Program.
  - Running the RSPFPAR program in a batch job named RUN\_PROGRAM\_RSPFPAR creates a value of RSPFPAR for Transaction code or program.

This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.

- The type is string. This is a key dimension.

transPerf Sample Interval Start

- The time stamp for the beginning time of the data supplied by the SAP agent.
- The type is timestamp.

transPerf Dialog Step Response Threshold (ms)



- The response time threshold, in milliseconds, for dialog steps. A dialog step with a response time that exceeds this threshold is counted in the Dialog Steps Above Threshold attribute. This value is set by configuring the SAP agent ABAP code. -1 = Not\_Set -2 = N/A.
- The type is int.

#### transPerf Executed in

- How the unit of work was started. For example
  - Running the FB01 transaction results in a u00e2u20acu201d value for Executed in to indicate a standalone transaction.
  - Running the RSPFPAR program through transaction SE38 creates a value of SE38 for Executed in.
  - Running the RSPFPAR program in a batch job named RUN\_PROGRAM\_RSPFPAR creates a value of RUN\_PROGRAM\_RSPFPAR for Executed in.
- The type is string. This is a key dimension.

#### transPerf SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### transPerf Row Aggregation

- Identifies the aggregation level that was used to create the row of data. See the Aggregation attribute above for all possible values. This attribute is most useful when viewing historical data records in the warehouse.
- The type is string.

#### transPerf Dynpro Number

- The Dynpro number referenced in the SAPGUI session.
- The type is string. This is a key dimension.

#### transPerf Sample Interval End

- The time stamp for the stopping time of the data supplied by the SAP agent.
- The type is timestamp.

#### transPerf Service Type

- The sap service type including: Dialog, Update, Batch and Spool.
- The type is string.

#### transPerf Logon Parameters

- This attribute is reserved for internal use only.
- The type is string.

#### transPerf System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### transPerf System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

#### transPerf Application (Unicode)

- The name of the business application name or of the sub-application name. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string. This is a key dimension.

#### transPerf Managed System

- The identifier for this SAP resource.
- The type is string. This is a key dimension.

#### transPerf Description

- The program name, transaction code, business application, or user ID description. This attribute provides single-byte character support only. Description is a language dependent description for the transaction code or program or how the unit of work was started. You control the type of reporting based on the value you set for the Aggregation attribute.
  - Examples of transaction performance data aggregated using the EXECIN aggregation level
    - Running the FB01 transaction creates a value of Standalone transaction for Description. (FB01 was executed as a transaction. )
    - Running the RSPFPAR program through transaction SE38 creates a value of ABAP Editor for Description. (ABAP Editor is obtained from the SAP system and is the language dependent description of the SE38 transaction. )
    - Running the RSPFPAR program in a batch job named RUN\_PROGRAM\_RSPFPAR creates a value of Job RUN\_PROGRAM\_RSPFPAR for Description. (This is the language dependent word for "Job" concatenated with the job name. )
  - Examples of transaction performance data aggregated without using the EXECIN aggregation level
    - Running the FB01 transaction creates a value of Post document for Description. (Post document is obtained from the SAP system and is the language dependent description of the FB01 transaction. )
    - Running the RSPFPAR program through transaction SE38 creates a value of Display Profile Parameter for Description. (Display Profile Parameter is obtained from the SAP system and is the language dependent title for the RSPFPAR program. )
    - Running the RSPFPAR program in a batch job named RUN\_PROGRAM\_RSPFPAR creates a value of Display Profile Parameter for Description. (Display Profile Parameter is obtained from the SAP system and is the language dependent title for the RSPFPAR program. )
- The type is string.

#### transPerf Program or Tran Code

- The unit of work that you started. It is determined by SAP code. For example
  - Running the FB01 transaction creates a value of FB01 for Transaction Code or Program.
  - Running the RSPFPAR program through transaction SE38 creates a value of RSPFPAR for Transaction code or Program.
  - Running the RSPFPAR program in a batch job named RUN\_PROGRAM\_RSPFPAR creates a value of RSPFPAR for Transaction code or program.

- The type is string.

transPerf SAPGUI Hostname

- Hostname of the SAPGUI logon terminal.
- The type is string.

transPerf Aggregation Level

- The aggregation level specifies how the transaction performance data is aggregated. This attribute is passed as input to the Transaction Performance data provider. The supports the following values: ALL = Deliver transaction performance for all top level aggregations, including the following aggregation levels TCODE, APPL, SUB, USERID, and EXECIN APPL = Report by Application name DYNPRO = Report by Program and Dynpro number EXECIN = Report by Executed In value that is the transaction or job that invoked the program. HIST = Row was collected as a result of historical data collection. This value is for internal use only. Do not set aggregation for this value. PEXEIN = Report by Transaction code or Program and Executed In value SUB = Report by Sub-application name TCODE = Report by Transaction code or program. TCUSER = Report by Transaction code or Program and user ID UEXEIN = Report by user Id and Executed In value USERID = Report by user ID All values except HIST are available for your use in workspace queries and situations. For workspace queries and situations, if no value is specified for the Aggregation attribute, the default value is ALL.
- The type is string.

transPerf Service Type Encoded

- The encoded sap service type.
- The type is string.

## Metrics

transPerf Max Extended Memory Per Session (kb) (Superseded)

- The maximum amount of extended memory, in KB, per session. For example, 132 indicates that the maximum amount of extended memory was 132 KB per session during the sampling period.
- The type is int.
- The unit is kilobytes.

transPerf Avg Response Time (ms) (Superseded)

- The average amount of time, in milliseconds, elapsed to process this transaction. For example, 177 indicates that the amount of time elapsed to process this transaction averaged 177 milliseconds during the sampling period.
- The type is int.
- The unit is milliseconds.

transPerf Avg Wait Time (ms)

- The average amount of time, in milliseconds, an unprocessed transaction waited in the queue for a free work process.
- The type is int.
- The unit is milliseconds.

transPerf Total Database Request Time (ms)

- The total amount of time, in milliseconds, elapsed for the database to process this transaction.
- The type is int.
- The unit is milliseconds.

transPerf Total DB Requested Bytes (kb)

- The total number of bytes, in KB, requested from the database for this transaction.
- The type is int.
- The unit is kilobytes.

transPerf Total Response Time (ms)

- The total amount of time, in milliseconds, elapsed to process this transaction.
- The type is int.
- The unit is milliseconds.

transPerf Avg Total Memory (kb) (Superseded)

- The average total amount of memory, in KB. For example, 5632 indicates that the total amount of memory is 5632 KB during the sampling period.
- The type is int.
- The unit is kilobytes.

transPerf GUI Time (ms)

- The number of milliseconds required to respond to a user SAPGUI request.
- The type is int.
- The unit is milliseconds.

transPerf GUI Time (ms) (Superseded)

- The number of milliseconds required to respond to a user SAPGUI request. This time is measured from when the user presses a key to send a request until the response is received.
- The type is int.
- The unit is milliseconds.

transPerf Dialog Steps Above Threshold

- Number of dialog steps with a response time that exceeded the threshold in the Dialog Step Response Threshold attribute.
- The type is int.
- The unit is steps.

transPerf Front End Network Time (ms) (Superseded)

- The number of milliseconds used in network communication. This is the GUI Time minus the application server processing time.
- The type is int.
- The unit is milliseconds.

transPerf Dialog Steps (Superseded)

- The number of dialog steps completed for this transaction. For example, 5 indicates that five dialog steps completed for this transaction during the sampling period.
- The type is int.
- The unit is steps.

transPerf Avg Database Request Time (ms)

- The average amount of time, in milliseconds, elapsed for the database to process this transaction.
- The type is int.
- The unit is milliseconds.

#### transPerf Max Extended Memory Per Session (kb)

- The maximum amount of extended memory, in KB, per session.
- The type is int.
- The unit is kilobytes.

#### transPerf Dialog Steps

- The number of dialog steps completed for this transaction.
- The type is int.
- The unit is steps.

#### transPerf Avg Database Request Time (ms) (Superseded)

- The average amount of time, in milliseconds, elapsed for the database to process this transaction. For example, 2 indicates that the amount of time elapsed to complete database requests for this transaction averaged 2 milliseconds during the sampling period.
- The type is int.
- The unit is milliseconds.

#### transPerf Avg Private Memory (kb)

- The average amount of private memory, in KB.
- The type is int.
- The unit is kilobytes.

#### transPerf Max Extended Memory Per Trans (kb) (Superseded)

- The maximum amount of extended memory, in KB, per transaction. For example, 2 indicates that the maximum amount of extended memory was 2 KB per transaction during the sampling period.
- The type is int.
- The unit is kilobytes.

#### transPerf Front End Network Time (ms)

- The number of milliseconds used in network communication.
- The type is int.
- The unit is milliseconds.

#### transPerf Avg Extended Memory (kb) (Superseded)

- The average amount of extended memory, in KB. For example, 132 indicates that the amount of extended memory averaged 132 KB during the sampling period.
- The type is int.
- The unit is kilobytes.

#### transPerf Total Response Time (ms) (Superseded)

- The total amount of time, in milliseconds, elapsed to process this transaction. For example, 333300 indicates that the amount of elapsed time, in milliseconds, to process this transaction totaled 333300 milliseconds during the sampling period.
- The type is int.
- The unit is milliseconds.

#### transPerf Avg Extended Memory (kb)

- The average amount of extended memory, in KB.
- The type is int.
- The unit is kilobytes.

#### transPerf Avg Wait Time (ms) (Superseded)

- The average amount of time, in milliseconds, an unprocessed transaction waited in the queue for a free work process. For example, 1 indicates that the amount of time an unprocessed transaction waited in the queue for a free work process averaged 1 millisecond during the sampling period.
- The type is int.
- The unit is milliseconds.

#### transPerf Total CPU Time (ms) (Superseded)

- The total amount of time, in milliseconds, that the CPU processed instructions for this transaction. For example, 180 indicates that the CPU processed instructions for this transaction for 180 milliseconds during the sampling period.
- The type is int.
- The unit is milliseconds.

#### transPerf Avg CPU Time (ms) (Superseded)

- The average amount of time, in milliseconds, the CPU processed instructions for this transaction. For example, 36 indicates that the amount of time the CPU processed instructions for this transaction averaged 36 milliseconds during the sampling period.
- The type is int.
- The unit is milliseconds.

#### transPerf Avg Total Memory (kb)

- The average total amount of memory, in KB.
- The type is int.
- The unit is kilobytes.

#### transPerf Total DB Requested Bytes (kb) (Superseded)

- The total number of bytes, in KB, requested from the database for this transaction. For example, 6144 indicates that a total of 6 MB were requested from the database for this transaction during the sampling period.
- The type is int.
- The unit is kilobytes.

#### transPerf Dialog Steps Above Threshold (%)

- Percentage of dialog steps with a response time that exceeds the threshold in the Dialog Step Response Threshold attribute. This attribute is calculated as Dialog Steps Above Threshold divided by Dialog Steps attribute.

- The type is int.
- The unit is percent.

#### transPerf Total Wait Time (ms)

- The total amount of time, in milliseconds, an unprocessed transaction waited in the queue for a free work process.
- The type is int.
- The unit is milliseconds.

#### transPerf Avg CPU Time (ms)

- The average amount of time, in milliseconds, the CPU processed instructions for this transaction.
- The type is int.
- The unit is milliseconds.

#### transPerf Max Extended Memory Per Trans (kb)

- The maximum amount of extended memory, in KB, per transaction.
- The type is int.
- The unit is kilobytes.

#### transPerf Dialog Steps Above Threshold (Superseded)

- Number of dialog steps with a response time that exceeded the threshold in the Dialog Step Response Threshold attribute.
- The type is int.
- The unit is steps.

#### transPerf GUI Count (Superseded)

- The number of roundtrip requests from a user workstation to the mySAP instance and back to the user workstation.
- The type is int.
- The unit is requests.

#### transPerf Total Database Calls (Superseded)

- The total number of database calls completed for this transaction. For example, 15 indicates that the application instance made a total of 15 requests to the database for this transaction during the sampling period.
- The type is int.
- The unit is calls.

#### transPerf Total Wait Time (ms) (Superseded)

- The total amount of time, in milliseconds, an unprocessed transaction waited in the queue for a free work process. For example, 2 indicates that the amount of time, in milliseconds, an unprocessed transaction waited in the queue for a free work process totaled 2 milliseconds during the sampling period.
- The type is int.
- The unit is milliseconds.

#### transPerf Avg Response Time (ms)

- The average amount of time, in milliseconds, elapsed to process this transaction.

- The type is int.
- The unit is milliseconds.

transPerf Total Database Request Time (ms) (Superseded)

- The total amount of time, in milliseconds, elapsed for the database to process this transaction. For example, 12 indicates that the amount of time elapsed to complete database requests for this transaction totaled 12 milliseconds during the sampling period.
- The type is int.
- The unit is milliseconds.

transPerf Total CPU Time (ms)

- The total amount of time, in milliseconds, that the CPU processed instructions for this transaction.
- The type is int.
- The unit is milliseconds.

transPerf Avg Private Memory (kb) (Superseded)

- The average amount of private memory, in KB. For example, 2612 indicates that the average amount of private memory is 2612 KB during the sampling period.
- The type is int.
- The unit is kilobytes.

transPerf Total Database Calls

- The total number of database calls completed for this transaction.
- The type is int.
- The unit is calls.

transPerf GUI Count

- The number of roundtrip requests from a user workstation to the SAP instance and back to the user workstation.
- The type is int.
- The unit is requests.

### **Component: Logon Information**

Logon Information is a system level data set that provides both current and historical information about users who have logged on to the mySAP system.

#### **Dimensions**

LogonInformation IP Address

- The IP address of the workstation being used. For example, 10. 20. 112. 14 is the IP address for the workstation.
- The type is string.

LogonInformation Userid Type

- The type of user ID.
- The type is string.

LogonInformation Logon Information User ID



- The name of the user logging on to the session. For example, RBROWN is the name of the user initiating the session.
- The type is string. This is a key dimension.

#### LogonInformation Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### LogonInformation Sample Interval End

- The time stamp for the stopping time of the data supplied by the SAP agent.
- The type is timestamp.

#### LogonInformation Changing Time

- The date and time when this user ID was locked or unlocked.
- The type is timestamp.

#### LogonInformation Logon Information Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### LogonInformation Logon Information Terminal

- A text string identifier or name for the computer terminal where the user logged on to the mySAP system. For example, LBROWN indicates the computer terminal.
- The type is string.

#### LogonInformation Logon Information System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

#### LogonInformation Logon Logoff

- The action presently occurring at the workstation. For example, Logon Pending indicates that a user is presently logging on to the workstation.
- The type is string.

#### LogonInformation Userid State

- The lock state of the user ID. The following values are possible: 0 = Not locked. User ID is currently not locked and there was no locking or unlocking activity on the user ID during the sample period. This user state is not reported by the ABAP unless the user ID has an invalid password count greater than 0. 1 = Locked. User ID is currently locked and there was no locking or unlocking activity on the user ID during the sample period. This user state is always reported. 2=Unlocked. User ID is currently not locked and was in a locked state at some time during the sample period. There was one or more unlocking activities on the user ID during the sample period with the last activity being an unlock. This user state is reported only during the sample period in which it is detected. 3=Relocked. User ID is currently locked and was in an unlocked state at some time during the sample period. There was one or more locking activities on the user ID during the sample period with the last activity being a lock. This user state is reported only during the sample period in which it is detected.
- The type is string.

#### LogonInformation SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### LogonInformation Instance Name

- The name of the application instance you are monitoring. For example, ddrum2\_PRD\_00 is the name of the application instance you are monitoring.
- The type is string.

#### LogonInformation Logon Information Time

- The time stamp for the date and time of the logon, the logoff, or the failed logon.
- The type is timestamp.

#### LogonInformation Sample Interval Start

- The time stamp for the beginning time of the data supplied by the SAP agent.
- The type is timestamp.

#### LogonInformation Changing UserID

- The user ID that locked or unlocked the user specified in the Userid attribute.
- The type is string.

#### LogonInformation IP Address (v4/v6)

- The IP address of the workstation being used. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

#### LogonInformation Client

- The name of the client to which you are logged on. For example, 800 is the name of the client to which you are logged on.
- The type is string.

#### LogonInformation System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

### Metrics

#### LogonInformation Invalid Password Count

- The current number of invalid logons for this particular user ID. For example, 3 indicates the current number of invalid logons for this particular user ID.
- The type is int.
- The unit is logons.

#### LogonInformation Session Duration (mins)

- The duration of the logon session, in minutes, calculated from the logon time and the logoff time. For example, 22 indicates the duration of the logon session, in minutes.
- The type is int.
- The unit is minutes.

## Component: SAP System Logs

System Log is an instance level data set that provides information about all messages written to the system log in a mySAP instance.

### Dimensions

#### Record Number

- The log record number.
- The type is string.

#### sapSystemLogs Terminal

- A text string identifier or name for the computer terminal where the user logged on to the mySAP system. For example, LBROWN indicates the computer terminal.
- The type is string.

#### sapSystemLogs Severity

- The number of critical system logs.
- The type is int.

#### sapSystemLogs System Name

- The SAP System Identifier (SID) for the mySAP system that you are monitoring. For example, PRD.
- The type is string.

#### Entry Time

- The time stamp for the date and time that the log entry was made.
- The type is timestamp. This is a key dimension.

#### sapSystemLogs Client

- A text string identifier or number for the originating client. Use this attribute to specify an identifier for a client. For example, 800 indicates the identifier for the originating client.
- The type is string.

#### Development class

- The identifier for the development class. For example, STUW is the identifier for the development class.
- The type is string.

#### sapSystemLogs SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### sapSystemLogs User

- A text string identifier or user ID for the user whose activities resulted in the log entry. For example, RSMITH indicates the user who generated the log entry.
- The type is string.

#### sapSystemLogs System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.

- The type is string.

#### Message Number

- A text string identifier or name for the system message. For example, S74 indicates the identifier for the system message.
- The type is string.

#### sapSystemLogs Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### sapSystemLogs Sample Interval End

- The time stamp for the stopping time of the data supplied by the SAP agent.
- The type is timestamp.

#### sapSystemLogs Program Name

- A unique identifier or name for the ABAP program that was running. This attribute provides single-byte character support only. For example, SAPLY210 indicates the name of the ABAP program.
- The type is string.

#### sapSystemLogs Program Name (Unicode)

- A unique identifier or name for the ABAP program that was running. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### sapSystemLogs Message Text

- Descriptive text associated with the system message. This attribute provides single-byte character support only. For example, CONVERSATION ID 53659 indicates the text of the system message.
- The type is string.

#### sapSystemLogs Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### sapSystemLogs Task Type

- A text string identifier for the type of task associated with the entry. For example, RD indicates the type of task associated with the entry.
- The type is string.

#### sapSystemLogs Message Text (Unicode)

- Descriptive text associated with the system message. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### Message Class

- A text string identifier or name for the category of the message.

- The type is string.

sapSystemLogs Transaction Code

- A unique identifier for the transaction whose processing resulted in the log entry. For example, A309 indicates the identifier for the transaction.
- The type is string.

sapSystemLogs Instance Name

- The name of the application instance that you are monitoring. For example, ddrum2\_PRD\_00 is the name of the application instance that you are monitoring.
- The type is string.

sapSystemLogs Sample Interval Start

- The time stamp for the beginning time of the data supplied by the SAP agent.
- The type is timestamp.

### **Metrics**

Record Count

- Count of system messages of a certain category. Reserved for use in queries to count messages by varying criteria.
- The type is int.
- The unit is messages.

### **Component: CCMS Current State**

CCMS Current State is an instance level data set that shows current state information from CCMS in the SAP system.

#### **Dimensions**

CCMSCurrState TID Internal Handle

- Internal handle for TID, used to link from the current state view to the alert view.
- The type is string.

CCMSCurrState System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the data base server associated with the target SAP system.
- The type is string.

CCMSCurrState Monitoring Types Class

- Class for the monitoring type.
- The type is string.

CCMSCurrState Monitoring Types Short Name

- Short name of monitoring type.
- The type is string.

CCMSCurrState Monitoring Types Full Name

- Full name of the monitoring type.
- The type is string. This is a key dimension.

#### CCMSCurrState Current State

- Current status of MTE.
- The type is string.

#### CCMSCurrState MT Index

- Index of MT in Tree, used for the topology view.
- The type is int.

#### CCMSCurrState Monitoring Types Number

- Monitoring type number range.
- The type is string.

#### CCMSCurrState Monitor Object Name

- Name of the monitoring object.
- The type is string.

#### CCMSCurrState CCMS Current State Managed System

- The identifier for this SAP resource.
- The type is string. This is a key dimension.

#### CCMSCurrState Monitoring Types ID

- Unique Identifier for monitoring types.
- The type is string.

#### CCMSCurrState CCMS Current State System Name

- The SAP System Identifier (SID) for the SAP system you are monitoring. For example, PRD.
- The type is string.

#### CCMSCurrState Monitor Set

- CCMS Monitor set to which this alert belongs.
- The type is string.

#### CCMSCurrState Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

#### CCMSCurrState Monitoring Segment Name

- Name of the monitoring segment.
- The type is string.

#### CCMSCurrState Instance Name

- Name of the application instance that you are monitoring, for example, DDRUM2\_PRD\_00.
- The type is string.

#### CCMSCurrState Monitor

- CCMS Monitor to which this alert belongs.
- The type is string.

CCMSCurrState Parent MT Index

- Index of Parent of MT in Tree, used for topology view.
- The type is int.

CCMSCurrState CCMS Current State Number

- Used for counting MTE state in chart view, or used as a flag in topology view.
- The type is int. This is a key dimension.

CCMSCurrState SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

CCMSCurrState Monitoring Context Name

- Name of the monitoring context.
- The type is string.

CCMSCurrState Customization Group Name

- Name of the customization group.
- The type is string.

### **Metrics**

CCMSCurrState Last Value Change Time

- Last value change time stamp.
- The type is timestamp.
- The unit is unspecified.

CCMSCurrState Occurrence Time

- Alert time stamp.
- The type is timestamp.
- The unit is unspecified.

### **SAP Process Integration**

Information about SAP Process Integration nodes.

#### **Component: Component Monitoring**

Component Monitoring is an instance level data set that provides an overview of the status of the different monitoring components in PI/XI. The component monitoring information is provided in Runtime Workbench.

#### **Dimensions**

PIcomponentMonitoring User Name

- User Name for login.
- The type is string.

PIcomponentMonitoring SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string. This is a key dimension.

PIcomponentMonitoring Component Monitoring URL

- Component Monitoring URL that redirects you to Runtime Workbench.
- The type is string.

PIcomponentMonitoring System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

PIcomponentMonitoring System Name

- The SAP System Identifier (SID) for the mySAP system that you are monitoring.
- The type is string.

PIcomponentMonitoring PicomponentMonitoring Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

### **Component: Integration Engine Background Job**

Background Job is an instance level data set that provides monitoring information of PI/XI Background Jobs and Job Logs.

#### **Dimensions**

BackgroundJobs Job ID

- ID of the background Job.
- The type is string. This is a key dimension.

BackgroundJobs Message Class

- Class associated with the background Job.
- The type is string.

BackgroundJobs Timestamp

- Job start date and time.
- The type is timestamp.

BackgroundJobs System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the database server associated with the target SAP system.
- The type is string.

BackgroundJobs System Name

- The SAP System Identifier (SID) for the SAP system you are monitoring. For example, PRD.
- The type is string.

BackgroundJobs Job Name

- Name of the background job.
- The type is string.

BackgroundJobs Message Number



- Message Number associated with the background job.
- The type is int.

#### BackgroundJobs Message Text

- Message text uncoded, including the parameters inserted and the text.
- The type is string.

#### BackgroundJobs SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### BackgroundJobs Job Created By

- The name of the SAP user who scheduled a job or a job-step to run. This user does not have to be the user who authorized the job to run.
- The type is string.

#### BackgroundJobs Message Type

- Type of background job that is shown in the log.
- The type is string.

#### BackgroundJobs Job Status

- Status of the Job.
- The type is string.

#### BackgroundJobs Integration Engine Background Job Managed System

- The identifier for this SAP resource.
- The type is string. This is a key dimension.

### **Metrics**

#### BackgroundJobs Sample Time

- The date and time that the agent collected data from SAP.
- The type is timestamp.
- The unit is unspecified.

### **Component: SAP PI System Details**

SAP Process Integration System Details.

#### **Dimensions**

##### PI Value

- The dummy field for the Value column in portrait mode.
- The type is string.

##### PI Assigned Update Instance

- The name of the application server assigned to a specific update server. For example, Updinst\_SY1\_00 is the instance configured with the mySAP update service for this application instance.
- The type is string.

#### PI Operation Mode

- A text string identifier or name for the current operation mode of the system. For example, Private indicates the current operation mode of the system. This attribute provides single-byte character support only.
- The type is string.

#### PI Database Type

- Type of database.
- The type is string. This is a key dimension.

#### PI Description

- The dummy field for the Description column in portrait mode.
- The type is string.

#### PI System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### PI Update Service Configured

- A Yes/No switch to indicate if the update service is configured.
- The type is string.

#### PI Instance Status

- The status of this application instance, either running or not running.
- The type is string. This is a key dimension.

#### PI Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### PI Sample Time

- The time stamp for the date and time the agent collected the data from mySAP.
- The type is timestamp.

#### PI Enqueue Service Configured

- A Yes/No switch to indicate if the enqueue service is configured.
- The type is string.

#### PI Central Instance

- A Yes/No switch to indicate if the application server is the central instance. This attribute can be useful when tailoring a situation.
- The type is string.

#### PI Instance Op Mode State

- The state in which the instance is included in the current operation mode of this application server.
- The type is string.

#### PI System Description

- A user-provided description of this application instance as defined in the mySAP system transport table. This attribute provides single-byte character support only.
- The type is string. This is a key dimension.

#### PI Spool Service Configured

- A Yes/No switch to indicate if the spool service is configured.
- The type is string.

#### PI Instance Start Time

- The time stamp for the date and time the application instance started.
- The type is timestamp. This is a key dimension.

#### PI\_Instance Host Name

- The name of the physical system, without the domain, on which this application server resides. For example, Insthost is the name of the application instance you are monitoring.
- The type is string. This is a key dimension.

#### PI Batch Service Configured

- A Yes/No switch to indicate if the batch service is configured.
- The type is string.

#### PI Instance Host IP Address

- The IP address of the physical system on which the application instance resides. For example, 170. 106. 1. 11 is the IP address of the physical system on which the application instance you are monitoring resides.
- The type is string. This is a key dimension.

#### PI Database Host IP Address

- The IP address of the physical system on which the database instance resides. This value is the same for all instances of a mySAP system. For example, 170. 106. 1. 1 is the IP address for the database host in the mySAP system you are monitoring.
- The type is string. This is a key dimension.

#### PI Message Service Configured

- A Yes/No switch to indicate if the message server is configured.
- The type is string.

#### PI Configuration String

- The services mask, or string, for this application server. For example, DVEBMGS indicates that the following mySAP services are configured for this instance: D = Dialog V = Update (stands for Verbucher in German) E = Enqueue B = Background M = Message server G = SNA gateway S = Spool.
- The type is string.

#### PI Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string. This is a key dimension.

#### PI Instance Host IP Address (v4/v6)

- The IP address of the physical system on which the application instance resides. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

#### PI Gateway Service Configured

- A Yes/No switch to indicate if the gateway service is configured.
- The type is string.

#### PI System Number

- The number assigned to this application server instance. For example, 01 is the number of the mySAP instance you are monitoring.
- The type is string. This is a key dimension.

#### PI Central Instance Name

- The name of the central instance application server that is configured for this mySAP system.
- The type is string.

#### PI System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

#### PI Database Host IP Address (v4/v6)

- The IP address of the physical system on which the database instance resides. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

#### PI Database Release

- Release associated with the database.
- The type is string.

#### PI Instance Name

- The name of the application server.
- The type is string. This is a key dimension.

#### PI System Release

- The release number for the level of software installed on this application server. For example, 640 indicates the level of software installed in the SAP mySAP system you are monitoring.
- The type is string. This is a key dimension.

#### PI Dialog Service Configured

- A Yes/No switch to indicate if the dialog service is configured.
- The type is string.

#### PI Database Host Name

- The name of the host computer running the database instance of a system. For example, DBhost is the name of the database host in the mySAP system you are monitoring.
- The type is string. This is a key dimension.

#### PI Operation Mode (Unicode)

- A text string identifier or name for the current operation mode of the system. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### PI SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### PI System Description (Unicode)

- A user-provided description of this application server instance as defined in the mySAP system transport table. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### PI Update2 Service Configured

- A Yes/No switch to indicate if the Update2 service is configured.
- The type is string.

#### PI System Start Time

- The time stamp for the date and time the system started.
- The type is timestamp.

#### PI Instance Stop Time

- The time stamp for the date and time the application instance stopped.
- The type is timestamp. This is a key dimension.

#### PI Database Name

- The name of the database instance defined for this mySAP system. This name is frequently the same as the mySAP SID, and is the same for each instance of a mySAP system. For example, DB4 is the name of the physical system on which the database server resides in the mySAP system you are monitoring.
- The type is string.

### **Metrics**

#### PI Update2 Stopped Percent

- Percent of Update2 work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Update Queue Percent

- Percentage of the dispatcher queue allotted for Update that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Update Complete Percent

- Percent of Update work processes in the Complete state.

- The type is int.
- The unit is percent.

#### PI Update Queue

- The number of tasks in the dispatch queue waiting for an Update work process.
- The type is int.
- The unit is tasks.

#### PI Spool Running Percent

- Percent of Spool work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Registered Users

- Number of total registered users currently for this SAP system.
- The type is int.
- The unit is users.

#### PI Active Users

- The current number of users logged on to this application instance. For example, 47 indicates the number of users currently logged on to the instance you are monitoring.
- The type is int.
- The unit is users.

#### PI Spool Processes

- The number of spool processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Enqueue Queue Percent

- Percentage of the dispatcher queue allotted for Enqueue that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Dialog Queue Percent

- Percentage of the dispatcher queue allotted for Dialog that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Dialog Running Percent

- Percent of Dialog work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Update2 Queue

- The number of tasks in the dispatch queue waiting for an Update2 work process.
- The type is int.

- The unit is tasks.

#### PI Batch Waiting Percent

- Percent of Batch work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### PI Update Waiting Percent

- Percent of Update work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### PI Batch Complete Percent

- Percent of Batch work processes in the Complete state.
- The type is int.
- The unit is percent.

#### Instances ConnectionFailed

- The total number of instances that have lost connection in the system.
- The type is int.
- The unit is instances.

#### PI System Up Duration

- The amount of time, in minutes, that the system has been up. For example, 12 indicates that the system has been up for 12 minutes. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### PI Dialog Processes

- The number of dialog processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Update2 Complete Percent

- Percent of Update2 work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Database Port

- Database Port.
- The type is int.
- The unit is port.

#### PI Enqueue Processes

- Number of enqueue work processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Instances Down

- The total number of application instances that are down in this system. This values are only reported for instances defined in an operation mode profile. For example, 3 indicates that 3 of the instances you are monitoring are not running.
- The type is int.
- The unit is instances.

#### PI Update2 Running Percent

- Percent of Update2 work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Batch Job Queue

- The number of batch jobs in Ready state.
- The type is int.
- The unit is jobs.

#### PI Update2 Queue Percent

- Percentage of the dispatcher queue allotted for an Update2 that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Dialog Stopped Percent

- Percent of Dialog work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Total GUI Sessions

- The total number of non-APPC-TM GUI sessions.
- The type is int.
- The unit is GUI sessions.

#### Instances Passive

- The total number of instances that are in passive state in the system.
- The type is int.
- The unit is instances.

#### PI Update Processes

- The number of update processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Enqueue Waiting Percent

- Percent of Enqueue work processes in the Waiting state.
- The type is int.
- The unit is percent.



#### PI Batch Processes

- The number of batch processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Spool Queue Percent

- Percentage of the dispatcher queue allotted for Spool that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Update2 Processes

- Number of Update2 work processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Dialog Queue

- The number of tasks in the dispatch queue waiting for a Dialog work process.
- The type is int.
- The unit is tasks.

#### PI Enqueue Running Percent

- Percent of Enqueue work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Total External Sessions

- The total number of user sessions (GUI and RFC).
- The type is int.
- The unit is user sessions.

#### PI Update2 Waiting Percent

- Percent of Update2 work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### PI Spool Queue

- The number of tasks in the dispatch queue waiting for a Spool work process.
- The type is int.
- The unit is tasks.

#### PI Interactive Users

- Number of interactive (GUI) users currently for this server.
- The type is int.
- The unit is users.

#### PI Spool Waiting Percent

- Percent of Spool work processes in the Waiting state.

- The type is int.
- The unit is percent.

#### PI Total RFC Sessions

- The total number of RFC sessions.
- The type is int.
- The unit is RFC sessions.

#### PI NoWP Queue

- The number of tasks in the dispatch queue waiting to be processed by the dispatcher itself or some other system service.
- The type is int.
- The unit is tasks.

#### PI Enqueue Queue

- The number of tasks in the dispatch queue waiting for an Enqueue work process.
- The type is int.
- The unit is tasks.

#### PI\_Instance Down Duration

- The amount of time, in minutes, an application instance has been down. For example, 12 indicates that a particular instance has been down for 12 minutes. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### PI Dialog Complete Percent

- Percent of Dialog work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Batch Running Percent

- Percent of Batch work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Enqueue Complete Percent

- Percent of Enqueue work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Total Active Users

- Number of active users currently for this server. It includes RFC users and interactive users.
- The type is int.
- The unit is users.

#### PI Dialog Waiting Percent

- Percent of Dialog work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### PI Spool Stopped Percent

- Percent of Spool work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Enqueue Stopped Percent

- Percent of Enqueue work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Spool Complete Percent

- Percent of Spool work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Update Running Percent

- Percent of Update work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI\_Instance Up Duration

- The amount of time, in minutes, an application instance has been up in this system. For example, 12 indicates that a particular instance has been up for 12 minutes. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### PI Batch Stopped Percent

- Percent of Batch work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Update Stopped Percent

- Percent of Update work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Instances Running

- The total number of instances that are running in this system. For example, 15 indicates that 15 instances you are monitoring are running.
- The type is int.
- The unit is instances.

#### PI RFC Users

- Number of RFC users currently for this server.
- The type is int.
- The unit is users.

### **Component: Integration Engine Job Overview**

Integration Engine Job Overview attributes provide information on the background jobs, such as the status of the jobs and if they are successful.

#### **Dimensions**

JobOverview Timestamp

- The data and time that the job started.
- The type is timestamp.

JobOverview IntEngJobOverview Managed System

- The identifier for this SAP resource.
- The type is string. This is a key dimension.

JobOverview Job Name

- Name of the background job.
- The type is string.

JobOverview System Name

- The SAP System Identifier (SID) for the SAP system you are monitoring. For example, PRD.
- The type is string.

JobOverview Job Status

- Status of a job.
- The type is string.

JobOverview Type

- Type of Job.
- The type is string.

JobOverview SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

JobOverview System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the data base server associated with the target SAP system.
- The type is string.

JobOverview Sample Interval Start

- The time stamp for the starting time of the data that is supplied by the SAP agent.
- The type is timestamp.

JobOverview Sample Interval End

- The time stamp for the stopping time of the data that is supplied by the SAP agent.

- The type is timestamp.

### **Metrics**

JobOverview Sample Time

- The time stamp for the date and time that the agent collected data from SAP.
- The type is timestamp.
- The unit is unspecified.

### **Component: BPE Monitoring**

Provides information about the XML message packaging status in the business process engine.

#### **Dimensions**

PIbpeMonitoring Quality of Service

- The quality of the service that runs the pipeline.
- The type is string.

PIbpeMonitoring System Name

- The SAP System Identifier (SID) for the SAP system you are monitoring. For example, PRD.
- The type is string.

PIbpeMonitoring PibpeMonitoring SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

PIbpeMonitoring System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the database server associated with the target SAP system.
- The type is string.

PIbpeMonitoring Relation between Message and Process Instance

- Relation between the message and the process instance.
- The type is string.

PIbpeMonitoring Managed System

- The identifier for this SAP resource.
- The type is string. This is a key dimension.

PIbpeMonitoring Status

- The status of the errors that occur.
- The type is string.

PIbpeMonitoring Message ID

- The ID associated with the message.
- The type is string. This is a key dimension.

PIbpeMonitoring Queue Name

- The name of the queue.
- The type is string.

#### PIbpeMonitoring Queue Assignment

- Queue assignment in Inbound Processing.
- The type is string.

#### PIbpeMonitoring Sample Interval Start

- The time stamp for the starting time of the data that is supplied by the SAP agent.
- The type is timestamp.

#### PIbpeMonitoring Message Packaging Mode

- Message packaging of the XI message mode.
- The type is string.

#### PIbpeMonitoring Sample Interval End

- The time stamp for the stopping time of the data that is supplied by the SAP agent.
- The type is timestamp.

### **Metrics**

#### PIbpeMonitoring Retry Count

- The number of failed delivery attempts.
- The type is int.
- The unit is attempts.

#### PIbpeMonitoring Configuration Version

- The version of the configuration for the message packaging in the Business Process engine for inbound processing.
- The type is int.
- The unit is unspecified.

#### PIbpeMonitoring Maximum\_Memory\_Per\_Message\_Package

- Maximum Memory allocated per Message Package of an XI Message.
- The type is int.
- The unit is GB.

#### PIbpeMonitoring Maximum Wait Time

- Maximum Wait Time of the oldest message in the message package.
- The type is int.
- The unit is seconds.

#### PIbpeMonitoring Number of Queues

- Number of queues per process type.
- The type is int.
- The unit is queues per process type.

#### PIbpeMonitoring Received Timestamp

- The date and time stamp that the message was received in the Business Process Engine.
- The type is timestamp.

- The unit is unspecified.

PIbpeMonitoring Maximum\_Number\_Of\_Messages

- Maximum number of messages for each message package.
- The type is int.
- The unit is messages.

PIbpeMonitoring Sample Time

- The time stamp for the date and time that the agent collected data from SAP.
- The type is timestamp.
- The unit is unspecified.

### **Component: Workflow Trace Logs**

Workflow trace logs is application level data set that provides information about workflow trace logs that occur in a mySAP PI/XI system. The workflow trace logs the important internal process flow information.

#### **Dimensions**

PIwfTrace Work Flow Trace Status

- Shows the status of the workflow trace.
- The type is string.

PIwfTrace Trace Level Description

- The description of the trace.
- The type is string.

PIwfTrace Trace Component

- Component to be traced.
- The type is string.

PIwfTrace Locally Visible

- This confirms if the instances of the workflow trace object have been created.
- The type is string.

PIwfTrace Trace Id

- Unique ID of the trace.
- The type is string. This is a key dimension.

PIwfTrace Parent Trace Id

- Trace ID of the parent workflow trace. A workflow might have sub-work and as a result a workflow trace might have a sub-workflow trace.
- The type is string.

PIwfTrace System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

PIwfTrace Work Flow Trace Logs Description

- Trace header descriptive text.
- The type is string.

#### PIwfTrace Work Flow Trace Logs System Name

- The SAP System Identifier (SID) for the mySAP system that you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

#### PIwfTrace Creator Name

- Name of the user who created the trace for the workflow. The valid format is an alphanumeric string, with a maximum of 15 characters.
- The type is string.

#### PIwfTrace Work Flow Trace Logs Managed System

- The identifier for this mySAP resource.
- The type is string.

#### PIwfTrace Trace Level

- Level of detail determining which trace entries are written. This show the numeric value for the trace.
- The type is string.

#### PIwfTrace Work Flow Trace Logs System

- Trace created by the system.
- The type is string.

### **Metrics**

#### PIwfTrace Creation Timestamp

- Date and time when the trace for the workflow was created.
- The type is timestamp.
- The unit is unspecified.

#### PIwfTrace Activated Timestamp

- Date and time when the trace for the workflow was activated.
- The type is timestamp.
- The unit is unspecified.

#### PIwfTrace Work Flow Trace Index Number

- The sequence number of the trace.
- The type is int.
- The unit is unspecified.

#### PIwfTrace Expiry Timestamp

- Date and time when the trace for the workflow expires.
- The type is timestamp.
- The unit is unspecified.

#### PIwfTrace Activation End Timestamp



- Date and time when the trace for the workflow ends.
- The type is timestamp.
- The unit is unspecified.

### **Component: Message Communication**

This data set monitors the communication between a sender capable of sending or receiving synchronous requests and a receiver capable of sending or receiving asynchronous responses.

#### **Dimensions**

PImessageCommunication SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

PImessageCommunication System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

PImessageCommunication Managed System

- The identifier for this mySAP resource.
- The type is string.

PImessageCommunication System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

#### **Metrics**

PImessageCommunication

- The time stamp for the date and time that the agent collected the data.
- The type is timestamp.
- The unit is unspecified.

### **Component: Persistent Layer Analysis**

Persistence layer analysis is an data set that provides information about the current configuration of the switch procedure and number of messages present in the SAP instance.

#### **Dimensions**

Messages Switch Mode

- Current mode of the switch.
- The type is string.

Messages Archived and Logically Deleted Messages

- Number of archived and logically deleted messages in the client.
- The type is int.

Messages Current Master Table

- Name of the current master table.
- The type is string.

#### Messages in VERS

- Number of messages in the VERS table.
- The type is int.

#### Messages in CLUP

- Number of messages in the CLUP table.
- The type is string.

#### Messages in Client

- Number of messages in the client.
- The type is int.

#### Messages in CLUR

- Number of messages in the CLURtable.
- The type is int.

#### Messages PIPersistentLayerAnalysis SAPshcut Parameters

- Parameters passed to saphscut for any transaction launch definition.
- The type is string.

#### Messages in MAST

- Number of messages in the MAST table.
- The type is int.

#### Messages to be Archived

- Number of messages to be archived in the client.
- The type is int.

#### Messages Instance Name

- Name of the instance for which you complete the persistence layer configuration.
- The type is string. This is a key dimension.

#### Messages in ERROR

- Number of messages in the ERROR table.
- The type is int.

#### Messages Current Fill Level in %

- Current fill level specified in percentage format.
- The type is string.

#### Messages in Database

- Number of messages in the database.
- The type is int.

#### Messages Current Container

- Name of the current container table.
- The type is string.

#### Messages PIPersistentLayerAnalysis Managed System

- The identifier for this SAP resource.
- The type is string.

#### Messages PIPersistentLayerAnalysis System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the data base server associated with the target SAP system.
- The type is string.

#### Messages in EMAST

- Number of messages in the EMAST table.
- The type is int.

#### Messages Logically Deleted Messages

- Number of logically deleted messages in the client.
- The type is int.

#### Messages PIPersistentLayerAnalysis System Name

- The SAP System Identifier (SID) for the SAP system you are monitoring. For example, PRD.
- The type is string.

#### Messages Reorganization Status

- Status of reorganization and if it required or not.
- The type is string.

#### Messages for Reorganization

- Number of messages for reorganization in the client.
- The type is int.

### **Metrics**

#### Messages Maximum Entries

- Maximum number of table entries for the master table.
- The type is int.
- The unit is entries.

#### Messages Number of Entries

- Number of table entries in the master table.
- The type is int.
- The unit is entries.

### **Component: Engine Status**

This data set shows the status of the business process engine component.

#### **Dimensions**

##### PIEngineStatus SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### PIEngineStatus System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID. DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### PIEngineStatus System Name

- SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

#### PIEngineStatus Class Name

- The object type name.
- The type is string.

#### PIEngineStatus Component

- Business Process Engine administration application name.
- The type is string. This is a key dimension.

#### PIEngineStatus SAP Server Current Time

- Current Time of the application server.
- The type is timestamp.

#### PIEngineStatus PiengineStatus Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### PIEngineStatus Process Type

- Business Process Engine administration type of process.
- The type is string. This is a key dimension.

#### PIEngineStatus Engine Status

- Status of the engine, for example, processing, running, error, or stop. The following values are available: S = Stopped P = In\_process R = Running E = Error.
- The type is string.

#### PIEngineStatus User Name

- PI/XI User Name.
- The type is string.

### **Component: XML Message Logs**

Provides information about the XML message.

#### **Dimensions**

##### PIxmlMsgLogs Message ID

- GUID for the Integration Engine objects.
- The type is string. This is a key dimension.

##### PIxmlMsgLogs Sending System

- Defines how an adapter transforms a message so that it is processed by the Integration Engine during inbound processing.
- The type is string.

#### PIxmlMsgLogs Period End

- Date and time that shows the end of the period.
- The type is timestamp.

#### PIxmlMsgLogs XML Message Logs Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### PIxmlMsgLogs Pipeline ID

- Integration Engine Pipeline ID.
- The type is string.

#### PIxmlMsgLogs System Name

- The SAP System Identifier (SID) for the mySAP system that you are monitoring. For example, PRD.
- The type is string.

#### PIxmlMsgLogs Inbound Interface Namespace

- This contains the inbound interface.
- The type is string.

#### PIxmlMsgLogs Period Start

- Date and time that shows the start of the period.
- The type is timestamp.

#### PIxmlMsgLogs Receiving System

- Defines how an adapter transforms a message so that it is processed by the receiver during outbound processing.
- The type is string.

#### PIxmlMsgLogs User Name

- User name of the SAP system.
- The type is string.

#### PIxmlMsgLogs Inbound Interface Name

- The name of the receiver who accepts the XML Message.
- The type is string.

#### PIxmlMsgLogs System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the data base server associated with the target SAP system.
- The type is string.

#### PIxmlMsgLogs SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

PIxmlMsgLogs Outbound Interface Namespace

- Contains the Outbound Interface.
- The type is string.

PIxmlMsgLogs Outbound Interface Name

- The name of the sender who sends the XML message.
- The type is string.

PIxmlMsgLogs Message Type

- Integration Engine Message type.
- The type is string.

### **Metrics**

PIxmlMsgLogs Initial Timestamp

- The initial time of the XML Message.
- The type is timestamp.
- The unit is unspecified.

PIxmlMsgLogs Execution From

- Time stamp that represent the execution date of XML message The valid format is time stamp.
- The type is timestamp.
- The unit is unspecified.

PIxmlMsgLogs Send Timestamp

- Date and time that the XML message was sent.
- The type is timestamp.
- The unit is unspecified.

### **SAP Solution Manager**

Information about SAP Solution Manager resources.

#### **Component: SAP Soution Manager System Details**

Provides infromation about SAP Solution Manager System.

#### **Dimensions**

PI Database Type

- Type of database.
- The type is string. This is a key dimension.

PI Value

- The dummy field for the Value column in portrait mode.
- The type is string.

PI Operation Mode (Unicode)

- A text string identifier or name for the current operation mode of the system. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### PI Update2 Service Configured

- A Yes/No switch to indicate if the Update2 service is configured.
- The type is string.

#### PI Configuration String

- The services mask, or string, for this application server. For example, DVEBMGS indicates that the following mySAP services are configured for this instance: D = Dialog V = Update (stands for Verbucher in German) E = Enqueue B = Background M = Message server G = SNA gateway S = Spool.
- The type is string.

#### PI Instance Op Mode State

- The state in which the instance is included in the current operation mode of this application server.
- The type is string.

#### PI Message Service Configured

- A Yes/No switch to indicate if the message server is configured.
- The type is string.

#### PI Instance Status

- The status of this application instance, either running or not running.
- The type is string. This is a key dimension.

#### PI Instance Host IP Address (v4/v6)

- The IP address of the physical system on which the application instance resides. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

#### PI Instance Stop Time

- The time stamp for the date and time the application instance stopped.
- The type is timestamp. This is a key dimension.

#### PI Batch Service Configured

- A Yes/No switch to indicate if the batch service is configured.
- The type is string.

#### PI Assigned Update Instance

- The name of the application server assigned to a specific update server. For example, Updinst\_SY1\_00 is the instance configured with the mySAP update service for this application instance.
- The type is string.

#### PI Instance Name

- The name of the application server.
- The type is string. This is a key dimension.

#### PI Operation Mode

- A text string identifier or name for the current operation mode of the system. For example, Private indicates the current operation mode of the system. This attribute provides single-byte character support only.
- The type is string.

#### PI Database Name

- The name of the database instance defined for this mySAP system. This name is frequently the same as the mySAP SID, and is the same for each instance of a mySAP system. For example, DB4 is the name of the physical system on which the database server resides in the mySAP system you are monitoring.
- The type is string.

#### PI Update Service Configured

- A Yes/No switch to indicate if the update service is configured.
- The type is string.

#### PI Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### PI Instance Host IP Address

- The IP address of the physical system on which the application instance resides. For example, 170. 106. 1. 11 is the IP address of the physical system on which the application instance you are monitoring resides.
- The type is string. This is a key dimension.

#### PI SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### PI Central Instance Name

- The name of the central instance application server that is configured for this mySAP system.
- The type is string.

#### PI System Start Time

- The time stamp for the date and time the system started.
- The type is timestamp.

#### PI Database Release

- Release associated with the database.
- The type is string.

#### PI Gateway Service Configured

- A Yes/No switch to indicate if the gateway service is configured.
- The type is string.



#### PI Central Instance

- A Yes/No switch to indicate if the application server is the central instance. This attribute can be useful when tailoring a situation.
- The type is string.

#### PI Enqueue Service Configured

- A Yes/No switch to indicate if the enqueue service is configured.
- The type is string.

#### PI Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string. This is a key dimension.

#### PI System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

#### PI Sample Time

- The time stamp for the date and time the agent collected the data from mySAP.
- The type is timestamp.

#### PI System Description (Unicode)

- A user-provided description of this application server instance as defined in the mySAP system transport table. This attribute provides multi-byte character support. Use this attribute for new queries and situations to obtain worldwide language support in your environment.
- The type is string.

#### PI Database Host Name

- The name of the host computer running the database instance of a system. For example, DBhost is the name of the database host in the mySAP system you are monitoring.
- The type is string. This is a key dimension.

#### PI System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

#### PI Database Host IP Address

- The IP address of the physical system on which the database instance resides. This value is the same for all instances of a mySAP system. For example, 170. 106. 1. 1 is the IP address for the database host in the mySAP system you are monitoring.
- The type is string. This is a key dimension.

#### PI Spool Service Configured

- A Yes/No switch to indicate if the spool service is configured.
- The type is string.

#### PI\_Instance Host Name

- The name of the physical system, without the domain, on which this application server resides. For example, Insthost is the name of the application instance you are monitoring.
- The type is string. This is a key dimension.

#### PI System Description

- A user-provided description of this application instance as defined in the mySAP system transport table. This attribute provides single-byte character support only.
- The type is string. This is a key dimension.

#### PI Instance Start Time

- The time stamp for the date and time the application instance started.
- The type is timestamp. This is a key dimension.

#### PI Database Host IP Address (v4/v6)

- The IP address of the physical system on which the database instance resides. This attribute is long enough to hold IPv4 or IPv6 addresses.
- The type is string.

#### PI Description

- The dummy field for the Description column in portrait mode.
- The type is string.

#### PI Dialog Service Configured

- A Yes/No switch to indicate if the dialog service is configured.
- The type is string.

#### PI System Number

- The number assigned to this application server instance. For example, 01 is the number of the mySAP instance you are monitoring.
- The type is string. This is a key dimension.

#### PI System Release

- The release number for the level of software installed on this application server. For example, 640 indicates the level of software installed in the SAP mySAP system you are monitoring.
- The type is string. This is a key dimension.

### Metrics

#### PI Instances Down

- The total number of application instances that are down in this system. This values are only reported for instances defined in an operation mode profile. For example, 3 indicates that 3 of the instances you are monitoring are not running.
- The type is int.
- The unit is instances.

#### PI Spool Queue

- The number of tasks in the dispatch queue waiting for a Spool work process.
- The type is int.
- The unit is tasks.

#### PI Spool Stopped Percent

- Percent of Spool work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI System Up Duration

- The amount of time, in minutes, that the system has been up. For example, 12 indicates that the system has been up for 12 minutes. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### PI Total Active Users

- Number of active users currently for this server. It includes RFC users and interactive users.
- The type is int.
- The unit is users.

#### PI Total External Sessions

- The total number of user sessions (GUI and RFC).
- The type is int.
- The unit is user sessions.

#### PI Update2 Waiting Percent

- Percent of Update2 work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### PI Spool Queue Percent

- Percentage of the dispatcher queue allotted for Spool that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Registered Users

- Number of total registered users currently for this SAP system.
- The type is int.
- The unit is users.

#### PI Batch Waiting Percent

- Percent of Batch work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### PI Dialog Queue

- The number of tasks in the dispatch queue waiting for a Dialog work process.
- The type is int.
- The unit is tasks.

#### PI Total RFC Sessions

- The total number of RFC sessions.
- The type is int.
- The unit is RFC sessions.

#### PI\_Instance Up Duration

- The amount of time, in minutes, an application instance has been up in this system. For example, 12 indicates that a particular instance has been up for 12 minutes. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### PI Enqueue Complete Percent

- Percent of Enqueue work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Update2 Complete Percent

- Percent of Update2 work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Enqueue Processes

- Number of enqueue work processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Spool Running Percent

- Percent of Spool work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Total GUI Sessions

- The total number of non-APPC-TM GUI sessions.
- The type is int.
- The unit is GUI sessions.

#### PI Dialog Processes

- The number of dialog processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Enqueue Queue

- The number of tasks in the dispatch queue waiting for an Enqueue work process.
- The type is int.
- The unit is tasks.

#### PI Update2 Processes

- Number of Update2 work processes running on this application instance.
- The type is int.
- The unit is processes.

#### Instances Passive

- The total number of instances that are in passive state in the system.
- The type is int.
- The unit is instances.

#### PI Update2 Queue Percent

- Percentage of the dispatcher queue allotted for an Update2 that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Update Stopped Percent

- Percent of Update work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Batch Processes

- The number of batch processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Update Running Percent

- Percent of Update work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Dialog Waiting Percent

- Percent of Dialog work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### PI Spool Complete Percent

- Percent of Spool work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Batch Stopped Percent

- Percent of Batch work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Update Waiting Percent

- Percent of Update work processes in the Waiting state.
- The type is int.

- The unit is percent.

#### PI Batch Complete Percent

- Percent of Batch work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Update Complete Percent

- Percent of Update work processes in the Complete state.
- The type is int.
- The unit is percent.

#### PI Enqueue Queue Percent

- Percentage of the dispatcher queue allotted for Enqueue that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI NoWP Queue

- The number of tasks in the dispatch queue waiting to be processed by the dispatcher itself or some other system service.
- The type is int.
- The unit is tasks.

#### PI Instances Running

- The total number of instances that are running in this system. For example, 15 indicates that 15 instances you are monitoring are running.
- The type is int.
- The unit is instances.

#### PI Batch Job Queue

- The number of batch jobs in Ready state.
- The type is int.
- The unit is jobs.

#### PI\_Instance Down Duration

- The amount of time, in minutes, an application instance has been down. For example, 12 indicates that a particular instance has been down for 12 minutes. A value of -1 indicates that there is no data at this time.
- The type is int.
- The unit is minutes.

#### PI Dialog Running Percent

- Percent of Dialog work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Update2 Running Percent

- Percent of Update2 work processes in the Running state.

- The type is int.
- The unit is percent.

#### PI Update Processes

- The number of update processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Enqueue Waiting Percent

- Percent of Enqueue work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### Instances ConnectionFailed

- The total number of instances that have lost connection in the system.
- The type is int.
- The unit is instances.

#### PI Update Queue Percent

- Percentage of the dispatcher queue allotted for Update that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Update Queue

- The number of tasks in the dispatch queue waiting for an Update work process.
- The type is int.
- The unit is tasks.

#### PI Update2 Stopped Percent

- Percent of Update2 work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Spool Processes

- The number of spool processes running on this application instance.
- The type is int.
- The unit is processes.

#### PI Database Port

- Database Port.
- The type is int.
- The unit is port.

#### PI Update2 Queue

- The number of tasks in the dispatch queue waiting for an Update2 work process.
- The type is int.
- The unit is tasks.

#### PI RFC Users

- Number of RFC users currently for this server.
- The type is int.
- The unit is users.

#### PI Interactive Users

- Number of interactive (GUI) users currently for this server.
- The type is int.
- The unit is users.

#### PI Enqueue Stopped Percent

- Percent of Enqueue work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Batch Running Percent

- Percent of Batch work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Dialog Stopped Percent

- Percent of Dialog work processes in the Stopped state.
- The type is int.
- The unit is percent.

#### PI Active Users

- The current number of users logged on to this application instance. For example, 47 indicates the number of users currently logged on to the instance you are monitoring.
- The type is int.
- The unit is users.

#### PI Spool Waiting Percent

- Percent of Spool work processes in the Waiting state.
- The type is int.
- The unit is percent.

#### PI Dialog Queue Percent

- Percentage of the dispatcher queue allotted for Dialog that is being used by waiting tasks.
- The type is int.
- The unit is percent.

#### PI Enqueue Running Percent

- Percent of Enqueue work processes in the Running state.
- The type is int.
- The unit is percent.

#### PI Dialog Complete Percent



- Percent of Dialog work processes in the Complete state.
- The type is int.
- The unit is percent.

### **Component: SAP BPM Alerts**

Provides information about the most important Business Processes in a SAP Solution Manager system.

#### **Dimensions**

sapbpmAlerts Client

- The number of the client.
- The type is string.

sapbpmAlerts Monitoring Type

- Monitoring type of business process alert.
- The type is string.

sapbpmAlerts SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

sapbpmAlerts Solution Id

- Solution Id of the solution defined in Solution Manager.
- The type is string.

sapbpmAlerts System Id

- The SAP System Identifier (SID) for the SAP system you are monitoring.
- The type is string.

sapbpmAlerts Monitoring Id

- Monitoring ID of the business process alert.
- The type is string.

sapbpmAlerts System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

sapbpmAlerts Alert Type

- The type of alert, for example, start, delay, or duration.
- The type is string.

sapbpmAlerts System Label

- System label generated from SID\_DBhostname, where SID is the target SAP system ID and DBhostname is the host name of the data base server associated with the target SAP system.
- The type is string.

sapbpmAlerts Alert rating

- The rating of the alert.

- The type is int.

sapbpmAlerts Alert Message

- The message associated with the alert.
- The type is string.

sapbpmAlerts SAP System

- Name of the SAP system.
- The type is string.

sapbpmAlerts Sample Interval Start

- The time stamp for the starting time of the data that is supplied by the SAP agent.
- The type is timestamp.

sapbpmAlerts Sample Interval End

- The time stamp for the stopping time of the data that is supplied by the SAP agent.
- The type is timestamp.

sapbpmAlerts Managed System

- The identifier for this SAP resource.
- The type is string. This is a key dimension.

sapbpmAlerts Alert Timestamp

- The date and time associated with the alert.
- The type is timestamp.

## Metrics

sapbpmAlerts Sample Time

- The time stamp for the date and time when the agent collected data from SAP.
- The type is timestamp.
- The unit is unspecified.

## Component: SAP PI Components

Provides information about the status of the PI components that are configured in PI domains in Solution Manager.

### Dimensions

sapPiComponents Self test Rating

- Status of the self test of the component in relation to the PI domain.
- The type is int.

sapPiComponents Component Stack

- Determines if the component is related to the ABAP stack or the Java stack.
- The type is string.

sapPiComponents Availability Rating

- Availability of the component to the PI domain.
- The type is int.

#### sapPiComponents Self test Status

- Availability of the component for self test.
- The type is string.

#### sapPiComponents Self test Time

- The self test time of the component.
- The type is timestamp.

#### sapPiComponents Availability Time

- Availability time of the component.
- The type is timestamp.

#### sapPiComponents Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### sapPiComponents System Label

- System label that is generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the database server that is associated with the target mySAP system.
- The type is string.

#### Component Name

- Name of the component that is monitored by the PI domain.
- The type is string.

#### sapPiComponents PI Domain Description

- Description of the PI domain that is configured in Solution Manager.
- The type is string.

#### Component Type

- Type of component that is monitored by the PI domain.
- The type is string.

#### sapPiComponents System Name

- The SAP System Identifier (SID) for the mySAP system that you are monitoring. For example, PRD.
- The type is string. This is a key dimension.

#### sapPiComponents Available Status

- Availability of the component for monitoring purposes.
- The type is string.

#### sapPiComponents SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### sapPiComponents PI Domain Name

- Name of the PI domain that is configured in Solution Manager. You might have more than one PI domain that is configured in Solution Manager.
- The type is string.

### **Metrics**

sapPiComponents Sample Time

- The time stamp for the date and time the agent collected the data.
- The type is timestamp.
- The unit is unspecified.

### **Component: SAP PI Channel Monitoring**

It is the Solution Level (SLM) level data set that provides information in relation to the PI channels in the PI domain that are configured in the Solution Manager work center.

### **Dimensions**

sapPIChannelMonitoring Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

PI Channel Name

- Name of the communication channel.
- The type is string.

sapPIChannelMonitoringDirection

- Specifies the direction of the communication channel, for example, the sender or receiver channel.
- The type is string.

sapPIChannelMonitoring Communication Component

- Displays the communication component to which the communication channel is assigned.
- The type is string.

sapPIChannelMonitoringLog Detail

- Describes the status of the communication channel. If an error occurs, the type of error is displayed.
- The type is string.

sapPIChannelMonitoring System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

sapPIChannelMonitoring Adapter Type

- Displays the type of adapter that is selected and configured for the communication channel.
- The type is string.

sapPIChannelMonitoring Adapter Engine

- Name of the adapter engine that is configured for the communication channel.

- The type is string.

sapPIChannelMonitoring Include for Alerting

- Indicates if a message is valid for alerting purposes.
- The type is string.

Processing Status

- Status of the communication channel.
- The type is int.

sapPIChannelMonitoring System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

Adapter Namespace

- Displays the namespace that contains the type of adapter.
- The type is string.

sapPIChannelMonitoring SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

sapPIChannelMonitoring Party

- Displays the communication party that is specified in the communication channel, if it is available.
- The type is string.

## Metrics

sapPIChannelMonitoring Sample Time

- The time stamp for the date and time the agent collected data from mySAP.
- The type is timestamp.
- The unit is unspecified.

sapPIChannelMonitoring Count

- Count of the PI Channel processing status. This attribute is not available for situations and it is used only in the PI Channel Processing Status graph view.
- The type is int.
- The unit is status count.

sapPIChannelMonitoring Last Refreshed Timestamp

- The time stamp on which the channel is processed. Include for Alerting indicates whether the channel is valid for alerting purposes.
- The type is timestamp.
- The unit is unspecified.

## Component: SAP SLM MAI System Monitoring

Provides configuration status and availability status information about SAP systems.n.

## Dimensions

### Product Version

- Description of the product version that is installed on the SAP System.
- The type is string.

### sapslmMAISystemMonitoring System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

### sapslmMAISystemMonitoring Performance

- Status of the critical performance indicators in the SAP System.
- The type is int.

### sapslmMAISystemMonitoring System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

### sapslmMAISystemMonitoring Configuration

- Configuration status of the system including results of configuration setting checks and updates to the configuration settings.
- The type is int.

### sapslmMAISystemMonitoring Exception

- Information about the error messages in the SAP system. 0=Not Applicable 1=Ok 2=Warnings 3=Not\_Ok 4=Unknown.
- The type is int.

### sapslmMAISystemMonitoring Availability

- Availability of the system and instance.
- The type is int.

### Technical System Name

- Name of the technical SAP system.
- The type is string.

### sapslmMAISystemMonitoring Configuration Status

- Overall configuration status of the SAP system.
- The type is int.

### sapslmMAISystemMonitoring Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

### sapslmMAISystemMonitoring SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

sapslmMAISystemMonitoring System Type

- Type of SAP system, for example, ABAP or JAVA.
- The type is string.

### Metrics

sapslmMAISystemMonitoring Count

- Overall count of the SAP systems by availability and configuration status.
- The type is int.
- The unit is SAP systems.

sapslmMAISystemMonitoring Sample Time

- The time stamp for the date and time the agent collected the data.
- The type is timestamp.
- The unit is unspecified.

### Component: SAP SLM Alerts

This group of data monitor for MAI Alerts.

#### Dimensions

sapslmAlerts Alert Id

- A unique ID that contains 32 characters. This ID identifies the alert within the MAI. This ID is not displayed on a user interface, but it is used to call other operations in the alert.
- The type is string. This is a key dimension.

Managed Object Type SLM Alert

- Type of Managed Object (MO), for example, Technical System. The following types of managed objects are available: 0=Blank/Space/Unspecified, 1=Database, 2=Host Server, 3=Process Integration (PI/XI) Domain, 4=Technical Instance, 5=Technical Component, 6=Technical System, 7=Technical Scenario, 8=Connection.
- The type is string.

sapslmAlerts Managed Object Type Enum

- Type of Managed Object, for example, Database. The following values are possible: 0=Blank/Space/Unspecified 1=Database, 2=Host Server, 3=Process Integration (PI/XI) Domain, 4=Technical Instance, 5=Technical Component, 6=Technical System, 7=Technical Scenario.
- The type is int.

sapslmAlerts Logon Parameters

- Parameters passed to ksar3 for any Take Action definition.
- The type is string.

sapslmAlerts Default Period

- Frequency that the alert collection mechanism runs.
- The type is int.

sapslmAlerts Managed System

- The identifier for this mySAP resource.
- The type is string. This is a key dimension.

#### Managed Object Name SLM Alert

- Name of the Managed Object (MO) that is associated with an alert.
- The type is string.

#### sapslmAlerts Alert Description

- Description of the alert. If a description of the alert is not available, nothing is shown.
- The type is string.

#### sapslmAlerts SAPshcut Parameters

- Parameters passed to sapshcut for any transaction launch definition.
- The type is string.

#### sapslmAlerts Alert Name

- Name that is associated with the alert. The Alert Name is the short name that relates to the alert.
- The type is string.

#### sapslmAlerts Technical Name (MEA)

- A technical ID that is associated with an alert. The technical ID is not guaranteed to be unique for a specific alert. However, two semantically different alerts usually have different technical names.
- The type is string.

#### sapslmAlerts Alert Severity

- Severity of an alert.
- The type is int.

#### sapslmAlerts Category

- A category, for example, Availability, that is associated with an alert. Each alert belongs to a category.
- The type is string.

#### sapslmAlerts Alert Priority

- Defined by the severity that is assigned to the alert during configuration and the current rating of an alert. By default, alerts are sorted by priority.
- The type is int.

#### sapslmAlerts System Name

- The SAP System Identifier (SID) for the mySAP system you are monitoring. For example, PRD.
- The type is string.

#### sapslmAlerts Alert Text Value

- Represents the text value of a metric if a metric contains a text value. Values are not returned for objects, such as alerts or events, or non-text values. This value is useful for metrics that contain some Error/Warning text without any numeric values.
- The type is string.

#### sapslmAlerts Status

- Status of an alert, for example, Open.



- The type is string.

sapslmAlerts Alert Timestamp

- Date and time that the alert occurred in the Solution Manager System. The format of the time stamp is the UTC time stamp in short form (YYYYMMDDhhmmss).
- The type is timestamp.

sapslmAlerts Object Type

- Type of object, for example, Alert.
- The type is string.

sapslmAlerts Alert Rating

- A rating, for example, Unknown, Green, Yellow or Red that is associated with an alert.
- The type is int.

sapslmAlerts System Label

- System label generated from SID\_DBhostname, where SID is the target mySAP system ID and DBhostname is the host name of the data base server associated with the target mySAP system.
- The type is string.

### Metrics

sapslmAlertsSample Time

- The time stamp for the date and time the agent collected the data.
- The type is timestamp.
- The unit is unspecified.

## SAP HANA DATABASE agent metrics

The metrics for SAP HANA DATABASE agent resource types collect data for monitoring with IBM Cloud App Management. Every SAP HANA DATABASE agent resource type defines a set of dimensions and metrics. The descriptions provide such information as data type, dimension key, and metric unit.

### SAP Hana Database

Information about SAP Hana Database.

#### Dimensions

DB Host

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

Node Name

- Manged system name.
- The type is string. This is a key dimension.

DB Name

- The name of the database. For example, SYSTEMDB, HTB, and PT1.
- The type is string. This is a key dimension.

SYS ID

- The system identifier (SID) of the SAP HANA database.
- The type is string. This is a key dimension.

### **Component: SAP Hana Database Host Information**

SAP Hana Database Host Information.

#### **Dimensions**

sapHanaHostInfo DB Host

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

sapHanaHostInfo Start Time

- The time when the first service was started on the host. This value is updated when the system is restarted.
- The type is string. This is a key dimension.

Active Status

- The status of database instance on the host. The status can be active or inactive.
- The type is string. This is a key dimension.

System ID

- The system identifier (SID) of the SAP HANA database.
- The type is string. This is a key dimension.

Timestamp

- This is the local time when the data was collected.
- The type is timestamp.

OS Name

- The operating system of the SAP HANA system.
- The type is string. This is a key dimension.

Originenode

- An unique identifier for SAP Hana Resource.
- The type is string. This is a key dimension.

Host Ip Address

- The IP address of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

Instance Number

- The administrative unit that consists of the server software components.
- The type is string.

Build Version

- The version of the SAP HANA database that is installed on the system.
- The type is string. This is a key dimension.

Is Master

- Provides information whether the SAP HANA system host is a master or not. This attribute returns value as master or slave.
- The type is string. This is a key dimension.

Port

- The port number of the SAP HANA host system.
- The type is string. This is a key dimension.

### **Component: SAP Hana DB Locks**

Provides information about the wait count of accumulated locks for the available services since the database was started.

#### **Dimensions**

DB Locks Node

- An identifier for SAP Hana Database resource.
- The type is string. This is a key dimension.

#### **Metrics**

DB Locks Count

- The lock wait count for the available services since the database was started.
- The type is int.
- The unit is times.

### **Component: SAP Hana DB Connection**

Provides details of the database connections.

#### **Dimensions**

DB\_Connection\_Status

- The status of the connection. The connection status can be running or idle.
- The type is string.

DB Connection Node

- An identifier for SAP Hana Database resource.
- The type is string. This is a key dimension.

#### **Metrics**

DB Connection Status Count

- The count of status of the connection.
- The type is int.
- The unit is connections.

### **Component: SAP Hana Database Table Size Information**

Provides information about the row and column tables.

#### **Dimensions**

Table Size Table Type

- Provides information about the type of the table. This attribute returns value as row or column or hybrid.
- The type is string. This is a key dimension.

#### Table Size Is Column Table

- Provides information whether the table is a column table or not. This attribute returns value as true or false.
- The type is string.

#### Table Size Schema Name

- The name of the schema of the database objects.
- The type is string. This is a key dimension.

#### Table Size Info Node

- An identifier for SAP Hana Database resource.
- The type is string. This is a key dimension.

#### Table Size Table Name

- The table name of the database tables.
- The type is string. This is a key dimension.

#### Table Size Is Rejected

- Provides information whether the table is replicated. This attribute returns value as true or false.
- The type is string.

#### Table Size Is Partitioned

- Provides information whether the table is a partitioned or not. This attribute returns value as true or false.
- The type is string.

### **Metrics**

#### Table Size Record Count

- The number of records in the database table.
- The type is int.
- The unit is records.

#### Table Size

- The allocated memory size (in MB) for the fixed-size and variable-size parts.
- The type is double.
- The unit is MB.

### **SAP Hana Database Host**

Provides Information About SAP HANA Database Host.

### **Dimensions**

#### Node1

- An identifier for SAP Hana Resource.
- The type is string. This is a key dimension.

#### Database Name

- The name of the database. For example, SYSTEMDB, HTB, and PT1.
- The type is string. This is a key dimension.

#### Database Host

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

#### SystemID

- The system identifier (SID) of the SAP HANA database.
- The type is string. This is a key dimension.

### **Component: SAP Hana Transaction Details**

Provides information about all the transactions that are created by the users or SAP HANA database.

#### **Dimensions**

##### Transaction Details Transaction ID

- The ID of the transaction that is running on HANA database.
- The type is string. This is a key dimension.

##### Transaction Details Transaction Status

- The status of the transaction. One of the following values is possible: inactive, active, precommitted, aborting, partial\_aborting, and active\_prepare\_commit.
- The type is string. This is a key dimension.

##### Transaction Details Last Commit ID

- The last commit ID of the transaction.
- The type is string.

##### Transaction Details Lock Wait Time

- The total wait time of the lock that was accumulated in the transaction.
- The type is string.

##### Transaction Details End Time

- The time when the transaction was completed.
- The type is string.

##### Transaction Details Current Statement ID

- The current statement ID of the transaction.
- The type is string.

##### Transaction Details Start Time

- The time when the transaction was started.
- The type is string.

##### Transaction Details Host

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

#### Transaction Details Port

- The internal port number of the SAP HANA system.
- The type is string.

#### Transaction Details Transaction Type

- The type of transaction, for example, user, version, and garbage collection.
- The type is string.

#### Transaction Details Node

- An identifier for SAP Hana Resource.
- The type is string. This is a key dimension.

#### Transaction Details Connection ID

- The ID of the connection that triggered the system operation.
- The type is string. This is a key dimension.

### **Metrics**

#### Transaction Details Acquired Lock count

- The number of locks that are acquired in the transaction.
- The type is int.
- The unit is locks.

#### Transaction Details Active Statement Count

- The number of opened cursors in the transaction.
- The type is int.
- The unit is cursors.

#### Transaction Details Lock Wait Count

- The number of locks that are waiting for the transaction.
- The type is int.
- The unit is locks.

### **Component: SAP Hana Host Information**

SAP Hana Database Host Information.

#### **Dimensions**

##### SAP Hana Host

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

##### Start Time

- The time when the first service was started on the host. This value is updated when the system is restarted.
- The type is string. This is a key dimension.

##### Active Status

- The status of database instance on the host. The status can be active or inactive.

- The type is string. This is a key dimension.

#### System ID

- The system identifier (SID) of the SAP HANA database.
- The type is string. This is a key dimension.

#### Timestamp

- This is the local time when the data was collected.
- The type is timestamp.

#### OS Name

- The operating system of the SAP HANA system.
- The type is string. This is a key dimension.

#### Originenode

- An identifier for SAP Hana resource.
- The type is string. This is a key dimension.

#### Host IP Address

- The IP address of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

#### Instance Number

- The administrative unit that consists of the server software components.
- The type is string.

#### Build Version

- The version of the SAP HANA database that is installed on the system.
- The type is string.

#### Is Master

- Provides information whether the SAP HANA system host is a master or not. This attribute returns value as master or slave.
- The type is string.

#### Port

- The port number of the SAP HANA host system.
- The type is string. This is a key dimension.

### **Component: SAP Hana Database Host**

Provides information about the statements for which execution time was greater than the configured threshold. If the system is multitenant, then the attributes provide information about the tenant database.

#### **Dimensions**

##### Expensive Statements Statement Hash

- The unique identifier for an SQL string.
- The type is string. This is a key dimension.

#### Expensive Statements Node

- An identifier for SAP Hana Resource.
- The type is string. This is a key dimension.

#### Expensive Statements Database User

- The user name that is used to connect to the database.
- The type is string. This is a key dimension.

#### Expensive Statements Connection ID

- The ID of the connection that triggered the system operation.
- The type is string. This is a key dimension.

#### Expensive Statements Host

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

#### Expensive Statements Start Time

- The timestamp when the execution of statement started.
- The type is string. This is a key dimension.

#### Expensive Statements Statement String

- The SQL statement that runs for duration longer than the defined threshold.
- The type is string. This is a key dimension.

#### Expensive Statements Statement ID

- The statement ID of the transaction.
- The type is string. This is a key dimension.

### **Metrics**

#### Expensive Statements Duration

- The time (in seconds) that is required for executing the statement.
- The type is int.
- The unit is seconds.

#### Expensive Statements Statement Count

- The total number of expensive statements in the SAP HANA database.
- The type is int.
- The unit is statements.

### **Component: SAP Hana Blocked Transactions**

Provides information about the transactions that are waiting to acquire transaction locks held by another transaction, network, or disk. If the system is multitenant, then the attributes provide information about the tenant database.

### **Dimensions**

#### Blocked Transactions Waiting Object Type

- The type of the object on which the lock is placed.
- The type is string.



#### Blocked Transactions Blocked Connection ID

- The connection ID of the blocked transaction.
- The type is string. This is a key dimension.

#### Blocked Transactions Lock Owner Transaction ID

- The ID of the transaction that is holding the lock.
- The type is string.

#### Blocked Transactions Waiting Schema Name

- The name of the schema on which the lock is placed.
- The type is string.

#### Blocked Transactions Node

- An identifier for SAP Hana Resource.
- The type is string. This is a key dimension.

#### Blocked Transactions Lock Owner Connection ID

- The connection ID associated with the write transaction that is holding the lock.
- The type is string. This is a key dimension.

#### Blocked Transactions Blocked Time

- The time from when the transaction is blocked.
- The type is string.

#### Blocked Transactions Lock Type

- The type of lock that is held by the blocking transaction. The lock type can be record, object, and metadata.
- The type is string.

#### Blocked Transactions Blocked Transaction ID

- The ID of the blocked transaction.
- The type is string. This is a key dimension.

#### Blocked Transactions Host

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

#### Blocked Transactions Blocked Update Transaction ID

- The ID of the blocked update transaction.
- The type is string. This is a key dimension.

#### Blocked Transactions Lock Owner Update Transaction ID

- The ID of the update transaction that is holding the lock.
- The type is string.

#### Blocked Transactions Lock Mode

- The access level of transactions to the locked record, table, and database. The lock mode can be shared, exclusive, and intentional exclusive.
- The type is string.

#### Blocked Transactions Waiting Object Name

- The name of the object on which the lock is placed.
- The type is string.

### **Component: SAP Hana Idle Cursor Details**

Provides information about the long idle cursors in the database.

#### **Dimensions**

##### Idle Cursor Index

- The index on the rows of the result set that a particular snapshot of a data collection represents. It is valid for SPS9 and SPS10.
- The type is string.

##### Idle Cursor Statement String

- The SQL statement that is currently running.
- The type is string.

##### Idle Cursor Node

- An Identifier for SAP Hana Resource.
- The type is string. This is a key dimension.

##### Idle Cursor Port

- The internal port number of the SAP HANA system.
- The type is string.

##### Idle Cursor Application User

- The application user that executes the cursor.
- The type is string. This is a key dimension.

##### Idle Cursor Snapshot ID

- The ID of the snapshot. The snapshot ID is valid for the SPS9 version.
- The type is string. This is a key dimension.

##### Idle Cursor Connection Status

- The status of the connection, such as running or idle.
- The type is string.

##### Idle Cursor Host

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

##### Idle Cursor Statement Status

- The status of SQL statement, such as none, active, or suspended.
- The type is string.

Idle Cursor Application Name

- The name of the application that executes the cursor.
- The type is string. This is a key dimension.

Idle Cursor Client Host

- The host name of client computer.
- The type is string.

### **Metrics**

Idle Cursor Timestamp

- This is the local time when the data was collected.
- The type is timestamp.
- The unit is unspecified.

Idle Cursor Idle Time

- The duration (in seconds) when the connection is unused and idle.
- The type is int.
- The unit is seconds.

### **Component: SAP Hana DB Cache Information**

Provides the aggregated information about the cache. If the system is multitenant, then the attributes provide information about the tenant database.

#### **Dimensions**

DB Cache Information Host Name

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

DB Cache Information Node

- An identifier for SAP Hana Resource.
- The type is string. This is a key dimension.

DB Cache Information Cache ID

- The ID for the cache instance.
- The type is string. This is a key dimension.

#### **Metrics**

DB Cache Information Miss Count

- The number of cache access attempts when the required data was not available in the cache instance.
- The type is int.
- The unit is misses.

DB Cache Information Insert Count

- The number of insertions that occurred in the cache instance.
- The type is int.
- The unit is inserts.

#### DB Cache Information Free Size Percent

- The percentage of total memory that is not used by the cache instance.
- The type is double.
- The unit is percent.

#### DB Cache Information Used Size

- The memory (in GB) that is used by the cache instance.
- The type is int.
- The unit is gigabytes.

#### DB Cache Information Total Size

- The maximum memory (in GB) that is available for the cache instance.
- The type is int.
- The unit is gigabytes.

#### DB Cache Information Entry Count

- The number of entries that are currently available in the cache instance.
- The type is int.
- The unit is entries.

#### DB Cache Information Invalidate Count

- The number of invalidations that occurred in the cache instance.
- The type is int.
- The unit is invalidations.

#### DB Cache Information Timestamp

- This is the local time when the data was collected.
- The type is timestamp.
- The unit is unspecified.

#### DB Cache Information Used Size Percent

- The percentage of memory that is used by the cache instance.
- The type is double.
- The unit is percent.

#### DB Cache Information Hit Count

- The number of cache access attempts when the required data was available in the cache instance.
- The type is int.
- The unit is hits.

#### DBCacheInformation Hit Count Percent

- The percentage of cache access attempts when the required data was available in the cache.
- The type is double.
- The unit is percent.

## **SAP Hana System**

Information about SAP HANA Database System.

### **Dimensions**

System\_Node

- An identifier for SAP Hana System resource.
- The type is string. This is a key dimension.

System\_ID

- The system identifier (SID) of the SAP HANA database.
- The type is string. This is a key dimension.

Build\_Version

- The version of the SAP HANA database that is installed on the system.
- The type is string.

Active\_Status

- The status of database instance on the host. The status can be active or inactive.
- The type is string.

System\_Port

- The port number of the SAP HANA host system.
- The type is string.

Is\_Master

- Provides information whether the SAP HANA system host is a master or not. This attribute returns value as master or slave.
- The type is string.

System\_Host

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

Host\_IP\_Address

- The IP address of the server that hosts the SAP HANA database.
- The type is string.

Start\_Time

- The time when the first service was started on the host. This value is updated when the system is restarted.
- The type is string.

Instance\_Number

- The administrative unit that consists of the server software components.
- The type is string.

OS\_Name

- The operating system of the SAP HANA system.

- The type is string.

### **Component: SAP Hana Resource Utilization**

Provides information about the usage of host resources, such as memory and processor.

#### **Dimensions**

Resource Utilization Node Name

- An identifier for SAP Hana System resource.
- The type is string. This is a key dimension.

Resource Utilization Host Name

- The name of the server that hosts the SAP HANA database.
- The type is string.

Resource Utilization Server Timestamp

- The server timestamp.
- The type is string.

Resource Utilization Snapshot ID

- The ID of the snapshot.
- The type is string. This is a key dimension.

#### **Metrics**

Resource Utilization Used Physical Memory

- The amount of physical memory that is used on the host system.
- The type is int.
- The unit is gigabytes.

Resource Utilization Wait Time Delta

- The amount of CPU time that is spent in waiting for I/O. On Linux systems, the actual value for the CPU waiting time is displayed. On Windows systems, this value is always displayed as 0.
- The type is int.
- The unit is seconds.

Resource Utilization User Time Delta

- The amount of CPU time that is spent the in user mode.
- The type is int.
- The unit is seconds.

Resource Utilization Total CPU System Time Delta

- The amount of CPU time that is spent in the kernel mode.
- The type is int.
- The unit is seconds.

Resource Utilization Used Swap Space

- The amount of swap memory that is used on the host system.
- The type is int.

- The unit is gigabytes.

#### Resource Utilization Free Physical Memory Percent

- The percentage of free physical memory on the host system.
- The type is int.
- The unit is percent.

#### Resource Utilization Used Physical Memory Percent

- The percentage of physical memory that is used on the host computer.
- The type is int.
- The unit is percent.

#### Resource Utilization CPU uUtilization Percent

- The percentage of CPU that is used by all the processes.
- The type is int.
- The unit is percent.

#### Resource Utilization Idle Time Delta

- The amount of time (in seconds) the CPU is not processing instructions.
- The type is int.
- The unit is seconds.

#### Resource Utilization Free Swap Space Percent

- The percentage of free swap memory on the host system.
- The type is int.
- The unit is percent.

### **Component: SAP Hana Host Alerts**

Provides information about the current alerts in the statistics server.

#### **Dimensions**

##### Host Alerts Node Name

- An identifier for SAP Hana System resource.
- The type is string. This is a key dimension.

##### Host Alert Details

- The information about the alert.
- The type is string.

##### Host Alert Name

- The short name of the alert check.
- The type is string.

##### Host Alerts Snapshot ID

- The ID of the snapshot. The time stamp when the alert was generated.
- The type is string.

##### Host Alerts Alert ID

- The ID of the current alert.
- The type is string. This is a key dimension.

#### Host Alert Rating

- The severity of the alert occurrence. The severity might be different for each alert. The following values are possible: 1 (information message), 2 (warning level 1), 3 (warning level 2), 4 (warning level 3), 5 (error message). Severity increases 1 - 5, 5 being the most critical.
- The type is int.

#### Host Alert Timestamp

- The time according to the local server time when the alert was occurred.
- The type is string.

#### Host Alert User Action

- The recommended steps to be performed by an administrator while checking the issues that are identified in an alert.
- The type is string.

### **Component: SAP Hana Expensive Statements**

Provides information about the statements for which execution time was greater than the configured threshold.

#### **Dimensions**

##### Expensive Statement Connection ID

- The ID of the connection that triggered the system operation.
- The type is string.

##### Expensive Statement Host

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

##### Expensive Statement Database User

- The user name that is used to connect to the database.
- The type is string.

##### Expensive Statement Start Time

- The timestamp when the execution of statement started.
- The type is string.

##### Expensive Statement Hash

- The unique identifier for an SQL string.
- The type is string.

##### Expensive Statement ID

- The statement ID of the transaction.
- The type is string.

##### Expensive Statement Node Name

- An identifier for SAP Hana System resource.



- The type is string.

#### Expensive Statement String

- The SQL statement that runs for duration longer than the defined threshold.
- The type is string.

### Metrics

#### Expensive Statement Duration

- The time (in seconds) that is required for executing the statement.
- The type is int.
- The unit is seconds.

#### Expensive Statement Count

- The total number of expensive statements in the SAP HANA database.
- The type is int.
- The unit is statements.

### Component: SAP Hana Disk Usage

Provides information about the disk configuration and usage of the host computer.

#### Dimensions

##### Disk Usage File System Type

- The type of file system on the SAP HANA host.
- The type is string.

##### Disk Usage Node Name

- An identifier for SAP Hana System resource.
- The type is string. This is a key dimension.

##### Disk Usage Device ID

- The internal ID of the database storage device.
- The type is string.

##### Disk Usage SubPath

- The name of the mount directory. For example, mnt00001.
- The type is string.

##### Disk Usage Host Name

- The name of the server that hosts the SAP HANA database. The host name is not set when the disk is used by more than one host.
- The type is string. This is a key dimension.

##### Disk Usage Type

- The type of disk usage on the host computer. The usage types are log, data, trace, data\_backup, and log\_backup.
- The type is string.

##### Disk Usage Path

- The path of the disk. For example, /hana/data/HTB.
- The type is string.

### **Metrics**

Diski Usage Used Size Percent

- The percentage of space usage of the disk.
- The type is int.
- The unit is percent.

Disk Usage UsedSize

- The space usage of the disk in GB.
- The type is int.
- The unit is gigabytes.

Disk Usage Total Size

- The maximum space usage of the disk in GB.
- The type is int.
- The unit is gigabytes.

### **Component: SAP Hana Garbage Collection**

Provides information about the garbage collection for the host computer.

#### **Dimensions**

Garbage Collection Port

- The internal port number of the SAP HANA system.
- The type is string. This is a key dimension.

Garbage Collection Node Name

- An identifier for SAP Hana System resource.
- The type is string. This is a key dimension.

Garbage Collection HostName

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

Garbage Collection Store Type

- The type of storage that is handled by the garbage collection job. The type of storage can be column store and livecache.
- The type is string.

Garbage Collection Volume ID

- The persistence volume ID.
- The type is string.

### **Metrics**

Garbage Collection Started Jobs

- The number of cleanup jobs that are started by the garbage collection process.

- The type is int.
- The unit is jobs.

#### Garbage Collection Queue Loads

- The number of all garbage collection queue loads.
- The type is int.
- The unit is queueLoads.

#### Garbage Collection Processed Jobs

- The number of undo files that are processed for cleanup by the garbage collector.
- The type is int.
- The unit is files.

#### Garbage Collection History Count

- The current count of history files in the garbage collection.
- The type is int.
- The unit is files.

#### Garbage Collection Waiter Count

- The current count of garbage collection waiters.
- The type is int.
- The unit is waiters.

### **Component: SAP Hana Database Details**

Provides information about the SAP HANA database.

#### **Dimensions**

##### Database Details Usage

- The usage of database for development, production, test, and customized operations.
- The type is string.

##### Database Details Database Status

- The status of the database, such as active, inactive, and unknown.
- The type is string.

##### Database Details Start Time

- The time when the database was started.
- The type is string.

##### Database Details Node

- An identifier for SAP Hana Resource.
- The type is string. This is a key dimension.

##### Database Details Database Name

- The name of the database. For example, SYSTEMDB, HTB, and PT1.
- The type is string. This is a key dimension.

##### Database Details Version

- The version of the database. The version format is major.minor.patch.build.
- The type is string.

#### Database Details System ID

- The system identifier (SID) of the SAP HANA database.
- The type is string.

#### Database Details Database Host

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

### **Component: SAP Hana Blocked Transactions**

Provides information about the transactions that are waiting to acquire transaction locks held by another transaction, network, or disk.

#### **Dimensions**

##### Blocked Transaction Connection ID

- The connection ID of the blocked transaction.
- The type is string.

##### Blocked Transaction Lock Owner Transaction ID

- The ID of the transaction that is holding the lock.
- The type is string.

##### Blocked Transaction Waiting Object Type

- The type of the object on which the lock is placed.
- The type is string.

##### Blocked Transaction Database Host

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

##### Blocked Transaction Lock Owner Connection ID

- The connection ID associated with the write transaction that is holding the lock.
- The type is string.

##### Blocked Transaction Blocked Time

- The time from when the transaction is blocked.
- The type is string.

##### Blocked Transaction Node Name

- An identifier for SAP Hana System resource.
- The type is string. This is a key dimension.

##### Blocked Transaction Update Transaction ID

- The ID of the blocked update transaction.
- The type is string.

Blocked TransactionLock Owner Update Transaction ID

- The ID of the update transaction that is holding the lock.
- The type is string.

Blocked Transaction Waiting Object Name

- The name of the object on which the lock is placed.
- The type is string.

Blocked Transaction ID

- The ID of the blocked transaction.
- The type is string.

Blocked Transaction Waiting Schema Name

- The name of the schema on which the lock is placed.
- The type is string.

Blocked Transaction Lock Type

- The type of lock that is held by the blocking transaction. The lock type can be record, object, and metadata.
- The type is string.

Blocked Transaction Lock Mode

- The access level of transactions to the locked record, table, and database. The lock mode can be shared, exclusive, and intentional exclusive.
- The type is string.

### **Component: SAP Hana System Information**

Provides information about the SAP HANA system.

#### **Dimensions**

System Info Node Name

- An identifier for SAP Hana System resource.
- The type is string. This is a key dimension.

System Info Instance Number

- The administrative unit that consists of the server software components.
- The type is string.

System Info Distributed System

- Provides information whether the SAP HANA system is a distributed system or not. This attribute returns value as yes or no.
- The type is string.

System Info Version

- The version of the SAP HANA database that is installed on the system.
- The type is string.

System Info Platform

- The operating system of the SAP HANA system.
- The type is string.

#### System Info Instance ID

- The instance ID of the SAP HANA database.
- The type is string.

### **Component: SAP Hana Service Status**

Provides information about the memory usage of the services in each host system.

#### **Dimensions**

##### Service Status Start Time

- The duration for which the service was running since the server was started.
- The type is string. This is a key dimension.

##### Service Status Service Name

- The name of the service.
- The type is string. This is a key dimension.

##### Service Status Node

- An identifier for SAP Hana System resource.
- The type is string. This is a key dimension.

##### Service Status Active Status

- The status of the service. The status can be no, yes, unknown, starting, and stopping.
- The type is string. This is a key dimension.

##### Service Status Host Name

- The name of the server that hosts the SAP HANA database.
- The type is string. This is a key dimension.

#### **Metrics**

##### Service Status Pending Requests

- The number of requests that are waiting in a queue to be processed by a service in the SAP HANA database server.
- The type is int.
- The unit is requests.

##### Service Status Available Memory

- The amount of memory that is available for use by the service.
- The type is double.
- The unit is gigabytes.

##### Service Status Thread Count

- The total number of threads for the service.
- The type is int.
- The unit is threads.

#### Service Status Response Time

- The time (in milliseconds) that is taken by the service to respond to requests from the clients.
- The type is double.
- The unit is milliseconds.

#### Service Status Open File Count

- The number of files that are currently opened through a service in the SAP HANA database server.
- The type is int.
- The unit is files.

#### Service Status Process Memory

- The amount of logical memory that is used by the service.
- The type is double.
- The unit is gigabytes.

#### Service Status Process CPU Time

- The total CPU usage (in seconds) by the current process since the service was started.
- The type is double.
- The unit is seconds.

#### Service Status Active Request Count

- The number of active requests that are being processed by the service.
- The type is int.
- The unit is requests.

#### Service Status Requests Per Second

- The number of requests that were issued per second to a service from the clients.
- The type is double.
- The unit is requests.

#### Service Status Active Thread count

- The number of threads that are actively processing requests for a service.
- The type is int.
- The unit is threads.

#### Service Status Process Physical Memory

- The amount of physical memory that is used by the service.
- The type is double.
- The unit is gigabytes.

#### Service Status Total CPU Time

- The total CPU usage (in seconds) by all the processes since the service was started.
- The type is double.
- The unit is seconds.

#### Service Status Total Memory

- The total amount of memory that is used by the service.
- The type is double.
- The unit is gigabytes.

### **Component: SAP Hana License Information**

Provides information about the currently valid SAP HANA license that is installed on the system.

#### **Dimensions**

##### SAP HANA License Expiration Date

- The expiration date of the validity period of the license.
- The type is string. This is a key dimension.

##### SAP HANA License Product Description

- The description of the licensed software product.
- The type is string.

##### SAP HANA License Installation Number

- The installation number of the SAP HANA database.
- The type is string.

##### SAP HANA License System Number

- The system number of the SAP HANA database.
- The type is string.

##### SAP HANA License Product Name

- The name of the licensed software product, such as SAP HANA.
- The type is string.

##### SAP HANA License Hardware Key

- The hardware key that is required for SAP HANA installation.
- The type is string. This is a key dimension.

##### SAP HANA License Node Name

- An identifier for SAP Hana System resource.
- The type is string. This is a key dimension.

##### SAP HANA License System ID

- The system identifier (SID) of the SAP HANA database.
- The type is string.

##### SAP HANA License Measurement Interval

- The interval of license measurement in hours.
- The type is string.

##### SAP HANA License First Installation Time

- The date and time of the first license installation.
- The type is string.



#### SAP HANA License Start Date

- The start date of the validity period of the license.
- The type is string.

#### SAP HANA License Locked Down

- Provides information whether the system is locked down due to license status or not. This attribute returns value as true or false.
- The type is string.

#### SAP HANA License Product Limit Description

- The description of the product usage to be measured and its unit.
- The type is string. This is a key dimension.

#### SAP HANA Licence Enforced

- Provides information whether the license has a product usage limit that is license type is enforced or not. This attribute returns value as true or false.
- The type is string.

#### SAP HANA License Last Successful Check

- The last known date when the license was successfully checked and found valid.
- The type is string.

#### SAP HANA License Permanent

- Provides information whether the license is a permanent or temporary. This attribute returns value as true or false.
- The type is string.

#### SAP HANA License Valid

- Provides information whether the license is valid or not. This attribute returns value as true or false.
- The type is string. This is a key dimension.

### **Metrics**

#### SAP HANA License Product Limit

- The duration of the product usage as specified in the license.
- The type is int.
- The unit is days.

#### SAP HANA License Product Usage

- The maximum value of product usage that occurred in the last 13 months. The product usage is measured periodically.
- The type is int.
- The unit is times.

### **Component: Statistic Server Alerts**

Statistic Server Alerts.

## Dimensions

### Server Alert Information

- The information about the alert.
- The type is string.

### Server Alert Generated Alert Time

- The ID of the snapshot. The time stamp when the alert was generated.
- The type is string.

### Server Alert Recommended Steps

- The recommended steps to be performed by an administrator while checking the issues that are identified in an alert.
- The type is string.

### Server Alert Id

- The ID of the current alert.
- The type is string. This is a key dimension.

### Server Alert Short Name

- The short name of the alert check.
- The type is string.

### Server Alert Time

- The time according to the local server time when the alert was occurred.
- The type is string.

### Server Alert Node

- Node.
- The type is string. This is a key dimension.

## Metrics

### Server Alert Severity

- The severity of the alert occurrence. The severity might be different for each alert. The following values are possible: 1 (information message), 2 (warning level 1), 3 (warning level 2), 4 (warning level 3), 5 (error message). Severity increases 1 - 5, 5 being the most critical.
- The type is int.
- The unit is severity.

## Component: SAP Hana Network IO

Provides information about the network I/O for the host computer.

### Dimensions

#### Network IO Receive Duration

- The time (in seconds) that is spent in receiving the input.
- The type is int.

#### Network IO Sender Port

- The port number that is used by the sending service.

- The type is string. This is a key dimension.

Network IO Node Name

- An identifier for SAP Hana System resource.
- The type is string. This is a key dimension.

Network IO Receiver Port

- The port number that is used by the receiving service.
- The type is string. This is a key dimension.

Network IO Receiver Host

- The host name of the receiving service.
- The type is string.

Network IO Sender Host

- The host name of the sending service.
- The type is string. This is a key dimension.

### **Metrics**

Network IO Receive Size

- The amount of data (in MB) that is received by the receiving service.
- The type is int.
- The unit is megabytes.

Network IO Send Duration

- The time (in seconds) that is spent in sending the output.
- The type is int.
- The unit is seconds.

Network IO Request Count

- The number of requests that are processed by the host computer.
- The type is int.
- The unit is requests.

Network IO Send Size

- The amount of data (in MB) that is sent by the sending service.
- The type is int.
- The unit is megabytes.

### **Sybase Server agent metrics**

The metrics for Sybase Server agent resource types collect data for monitoring with IBM Cloud App Management. Every Sybase Server agent resource type defines a set of dimensions and metrics. The descriptions provide such information as data type, dimension key, and metric unit.

## **sybaseDatabase**

Sybase Database.

### **Dimensions**

Truncate Log on CKPT

- Indicates whether the database truncate log on checkpoint is enabled.
- The type is string.

Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

Sybase Database Detail Server

- The name of the Sybase server.
- The type is string.

Host Name

- The host name where the Sybase server resides.
- The type is string.

Select Into Bulkcopy Enabled

- Indicates whether the database is enabled for select into bulk copy.
- The type is string.

Database Name

- The name of the database.
- The type is string.

DB Owner (Unicode)

- The owner of the database.
- The type is string.

Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

Single User Access

- Indicates whether the database is set for single user only access.
- The type is string.

DBO Only Access

- Indicates whether the database is set for owner only access.
- The type is string.

Read Only Access

- Indicates whether database in read only mode.
- The type is string.

#### DB Owner

- The owner of the database.
- The type is string.

#### SybaseDatabaseDetail Row Number

- The row number in sample.
- The type is int.

#### Sybase Database Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### DB ID

- The ID of the database.
- The type is string. This is a key dimension.

#### Abort Tran On Log Full

- Indicates whether the abort tran on log full option is enabled for the database.
- The type is string.

#### Database Name (Unicode)

- The name of the database.
- The type is string.

### Metrics

#### Allocated Size In KB

- The amount of space in kilobytes allocated in the database.
- The type is int.
- The unit is kb.

#### Index Size In KB

- The amount of space in kilobytes used for the indexes in the database.
- The type is int.
- The unit is kb.

#### Data Size (MB)

- The number of megabytes (MB) allocated for the data only segments of the database.
- The type is double.
- The unit is mb.

#### Used Size In KB

- The amount of space in kilobytes used in the database.
- The type is double.
- The unit is kb.

#### Log Freespace (MB)

- The number of megabytes (MB) of free space of the transaction log in the database.

- The type is double.
- The unit is mb.

#### Log Freespace Percent

- The percentage of free space in the transaction log of the database.
- The type is double.
- The unit is percent.

#### Dump Tran Date

- The timestamp that indicates the date on which the dump transaction command was last executed for the database in the format yy.mm.dd.
- The type is timestamp.
- The unit is timestamp.

#### Data Freespace Percent

- The percentage of maximum available free space in the database.
- The type is double.
- The unit is percent.

#### Total Devices

- The number of devices allocated for the database.
- The type is int.
- The unit is devies.

#### Data Freespace (MB)

- The number of megabytes (MB) of free space in the database.
- The type is double.
- The unit is mb.

#### Free Space Accounting Suppresed

- Indicates whether the free space accounting option is suppresed for the database.
- The type is string.
- The unit is status.

#### Log Size (MB)

- The number of megabytes (MB) allocated for the transaction log of the database.
- The type is double.
- The unit is mb.

#### Error Status

- Indicates whether the database has an error status.
- The type is string.
- The unit is status.

#### Unallocated Size In KB

- The amount of space in kilobytes unallocated in the database.
- The type is double.

- The unit is kb.

#### No CKPT After Recovery

- Indicates whether a record for the checkpoint is added to the transaction log when the database is recovered.
- The type is string.
- The unit is status.

### **Component: resource**

It contains Sybase device details.

#### **Dimensions**

##### Server Version

- The version of the server.
- The type is string.

##### Physical Device Name

- The name of the physical device.
- The type is string.

##### SybaseDevice Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

##### Device Size (MB)

- The device size in megabytes.
- The type is string. This is a key dimension.

##### Sybase Device Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

##### Device Name. (Unicode)

- The device name.
- The type is string. This is a key dimension.

##### Device Type

- Indicates the type of device (data, log or data and log).
- The type is string.

##### Sybase Device Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

##### Physical Device Name (Unicode)

- The name of the physical device.
- The type is string.

#### Sybase Device Database Name

- The database name.
- The type is string.

#### Mirror Device Name (Unicode)

- The name of the mirror device.
- The type is string.

#### Mirror Device Name

- The name of the mirror device.
- The type is string.

#### SybaseDevice Row Number

- The row number in sample.
- The type is int.

#### Sybase Device Host Name

- The host name on which Sybase server resides.
- The type is string.

#### Device Name

- The device name.
- The type is string. This is a key dimension.

#### The name of the Sybase server.

- The name of the Sybase server.
- The type is string.

#### Database Name. (Unicode)

- The database name.
- The type is string.

### **Metrics**

#### Device Free Space (MB)

- The amount of free space in megabytes on this device .
- The type is double.
- The unit is mb.

#### Device Free Space Percent

- The percentage of free space on the device.
- The type is double.
- The unit is percent.

### **Component: Sybase Locks**

It contains Sybase\_Locks.

#### **Dimensions**

SybaseLocks CclassU



- The name of the cursor that is associated with the lock, if any.
- The type is string.

#### SybaseLocks Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### SybaseLocks Row Number

- The row number in sample.
- The type is int.

#### Sybase Locks Host Name

- The host name on which the server resides.
- The type is string.

#### Sybase Locks Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

#### Sybase Locks

- Indicates the types of lock: Exclusive Table, Shared Table, Exclusive Intent, Shared Intent, Exclusive Page, Shared Page, Update Page, Exclusive Extent, Update Extent, Next Extent, Previous Extent, Blocking Exclusive Table, Blocking Update Extent.
- The type is string. This is a key dimension.

#### SybaseLocks Server

- The name of the SQL server.
- The type is string. This is a key dimension.

#### Class

- The name of the cursor that is associated with the lock, if any.
- The type is string.

#### Sybase Locks Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

#### Sybase Locks Database Name (Unicode)

- The name of the database which is locked.
- The type is string.

#### Table Name (Unicode)

- The name of the table which has been locks.
- The type is string.

#### Sybase Locks Database Name

- The database name.
- The type is string. This is a key dimension.

#### Sybase Locks Database Id

- The ID of the database.
- The type is int.

#### **Metrics**

##### Table Name

- The name of table which has been locked.
- The type is string.
- The unit is name.

##### Page Number

- The page number of table which has been locked.
- The type is int.
- The unit is page.

##### Process Holding Lock

- The process ID of the lock holder.
- The type is int.
- The unit is id.

#### **sybaseInstance**

Sybase Instance.

#### **Dimensions**

##### Sample Timestamp

- The timestamp that indicates the date and time the product collected the sample for the server.
- The type is timestamp.

##### Sybase Server Detail Server

- The name of the Sybase server.
- The type is string.

##### Host Name

- The host name where the Sybase server resides.
- The type is string.

##### Server Version

- The version of the server.
- The type is string.

##### Procedure Cache Size (KB)

- The number of kilobytes (KB) that are allocated for the procedure cache.
- The type is int.

##### Sybase Server Detail FQDN

- The fully qualified host name where the Sybase server resides.
- The type is string.

#### Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

#### Port Number

- PORT NUMBER FOR SYBASE SERVER COMMUNICATION.
- The type is int.

#### Startup Timestamp

- The timestamp that indicates the date and time the server was started.
- The type is timestamp.

#### Error Log Name

- The name of the file that contains the error log for the server.
- The type is string.

#### Sybase Server Detail Row Number

- The row number in sample.
- The type is int.

#### Error Log Name (Unicode)

- The name of the file that contains the error log for the server.
- The type is string.

#### Max Locks Allowed

- The maximum number of locks permitted.
- The type is int.

#### OS Version

- The version of the operating system of the server.
- The type is string.

#### Sybase Server Detail Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### Server Type

- The type of Server. It indicates whether the server is SQL Server or backup server.
- The type is string.

#### OS Type

- The operating system of the server.
- The type is string.

#### Data Cache Size (KB)

- The number of kilobytes (KB) allocated for the data buffer cache.
- The type is int.

## Metrics

### Current Locks

- The number of current locks for the server.
- The type is int.
- The unit is locks.

### Backup Server Status

- Indicates the status of Backup Server, whether it is Active or Inactive.
- The type is string.
- The unit is status.

### Server Status

- Indicates the status of the Sybase server (Active, Inactive, Unknown).
- The type is string.
- The unit is status.

### Procedure Cache Percent

- The percentage of cache memory the server uses for the procedure cache.
- The type is double.
- The unit is percent.

### Error Log Size (Bytes)

- The number of bytes in the error log file.
- The type is int.
- The unit is bytes.

### Percent Max Locks

- The percentage of locks on resources of the maximum number of locks allowed by the server.
- The type is double.
- The unit is percent.

### Time Since Startup (Hrs.)

- The number of hours that have elapsed since the server was started.
- The type is int.
- The unit is hours.

### Job Server Status

- Indicates the status of Job Server, whether it is Active or Inactive.
- The type is string.
- The unit is status.

## Component: Sybase\_Statistics\_Summary

It contains Sybase\_Statistics\_Summary.

### Dimensions

Sybase Statistics Summary Sample Timestamp

- The time when these data were collected.

- The type is timestamp.

#### SybaseStatisticsSummary Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### Max User Connections Allowed

- The maximum permitted user connections on the server.
- The type is int.

#### Sybase Statistics Summary Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

#### Sybase Statistics Summary Current Interval (Sec.)

- The duration of current interval in seconds.
- The type is int.

#### Sybase Statistics Summary Value

- The attribute value.
- The type is string.

#### Sybase Statistics Summary \*

- The attribute name.
- The type is string.

#### Sybase Statistics Summary Host Name

- The host name where the Sybase server resides.
- The type is string.

#### Sybase Statistics Summary Row Number

- The row number in sample.
- The type is int.

#### Sybase Statistics Summary Server

- The name of the Sybase server.
- The type is string.

### **Metrics**

#### Disk IO Current Interval

- The number of disk IOs operations during the current interval.
- The type is int.
- The unit is count.

#### Total OS IO Pct Busy

- The percentage of IO busy during the current interval.
- The type is double.

- The unit is percent.

#### Physical Writes (per Sec.)

- The number of disk writes operations per second during the current interval.
- The type is double.
- The unit is count.

#### Total OS CPU Pct Busy

- The percentage of CPU busy during the current interval.
- The type is double.
- The unit is percent.

#### Pct IO Errors Cur Intvl

- The percentage of disk errors in relation to the total disk IO during the current interval.
- The type is double.
- The unit is percent.

#### Logons Available

- The number of logons or connections available.
- The type is int.
- The unit is count.

#### Physical Reads (per Sec.)

- The number of disk reads operations per second during the current interval.
- The type is double.
- The unit is count.

#### IO Errors Since Startup

- The number of disk errors since server startup.
- The type is int.
- The unit is count.

#### IO Errors Current Interval

- The number of disk errors during the current interval.
- The type is int.
- The unit is count.

#### Total Logons (per Sec.)

- The number of logon per second during the current interval.
- The type is double.
- The unit is count.

#### Current Logons

- The number of active logons or connections.
- The type is int.
- The unit is connections.

Pct Max Logons Active

- The percentage active logons or connections in relation to the maximum users.nconnections.
- The type is double.
- The unit is percent.

### **Component: Sybase\_Problem\_Summary**

It contains information about Sybase Problem Summary.

#### **Dimensions**

Sybase Problem Summary Current Interval (Sec.)

- The duration in seconds for the current interval.
- The type is int.

Sybase Problem Summary Row Number

- The row number in sample.
- The type is int.

Maximum Sev Current Interval

- The highest severity level of error messages during the current interval.
- The type is string.

Maximum Sev Timestamp

- The timestamp of the highest severity level of error messages since startup.
- The type is timestamp.

Sybase Problem Summary Host Name

- The host name where the Sybase server resides.
- The type is string.

Sybase Problem Summary Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

Maximum Sev Level

- The highest severity level of error messages since startup.
- The type is string.

Sybase Problem Summary Server

- The name of the Sybase server.
- The type is string.

Sybase Problem Summary Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

Sybase Problem Summary Age of Last Error (Min.)

- The age in minutes of the last error encountered.
- The type is int.

Sybase Problem Summary V alue

- The attribute value.
- The type is string.

Sybase Problem Summary Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

Sybase Problem Summary Description

- The attribute name.
- The type is string.

## Metrics

Total Errors High Sev

- The number of error messages of severity 17 and higher since startup.
- The type is int.
- The unit is count.

Total Error Messages

- The number of error messages since startup.
- The type is int.
- The unit is count.

Total Errors Cur Intvl

- The number of error messages during the current interval.
- The type is int.
- The unit is count.

Total Errors Other

- The number of Other error messages since startup.
- The type is int.
- The unit is count.

Sybase Problem Summary Error Log Size (Bytes)

- The size of error log in bytes.
- The type is int.
- The unit is bytes.

## Component: sybaseProblemDetail

It contains sybaseProblemDetail.

### Dimensions

Message Text

- The message text.
- The type is string. This is a key dimension.

Sybase Problem Detail Sample Timestamp



- The time when these data were collected.
- The type is timestamp.

#### Severity Level

- The severity of the error.
- The type is string. This is a key dimension.

#### Maximum Severity

- The maximum severity code.
- The type is string.

#### Message Issuer (Unicode)

- The message issuer.
- The type is string.

#### SybaseProblemDetail Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### SybaseProblemDetail Row Number

- The row number in sample.
- The type is int.

#### SybaseProblemDetail Server

- The name of the Sybase server.
- The type is string.

#### Message Timestamp

- The time when the the error message was issued.
- The type is timestamp. This is a key dimension.

#### Message Text (Unicode)

- The message text.
- The type is string.

#### Sybase Problem Detail Host Name

- The host name where the Sybase server resides.
- The type is string.

#### Error ID (Unicode)

- The error ID.
- The type is string.

#### Sybase Problem Detail Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

#### Error ID

- The error ID.
- The type is string. This is a key dimension.

#### Message Issuer

- The message issuer.
- The type is string.

### Metrics

#### SQL State Code

- The SQL state or the return code.
- The type is int.
- The unit is code.

#### Message Age (Min.)

- The age in minutes since message was issued.
- The type is int.
- The unit is minutes.

### Component: sybaseSQLDetail

It contains sybase SQL detail.

#### Dimensions

##### SybaseSQLDetail Server

- The name of the Sybase server.
- The type is string.

##### SQL Activity Name

- The name of the SQL activity (updated direct, inserted heap, deleted deferred).
- The type is string. This is a key dimension.

##### Sybase SQL Detail Host Name

- The host name where the Sybase server resides.
- The type is string.

##### Sybase SQL Detail Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

##### Sybase SQL Detail Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

##### SybaseSQLDetail Row Number

- The row number in sample.
- The type is int.

##### SybaseSQLDetail Originnode

- The key for the table view in the format SYB.serverid.hostname.

- The type is string. This is a key dimension.

## **Metrics**

### SQL Activity Count

- The count number of occurrences of SQL activity during the sampling period.
- The type is int.
- The unit is count.

### SQL Activity (per Trans.)

- The count number of occurrences of SQL activity per transaction.
- The type is double.
- The unit is activity per trans.

### SQL Activity Percent

- The percentage of activity represented by the specified activity.
- The type is double.
- The unit is percent.

### SQL Activity (per Sec.)

- The count number of occurrences of the SQL activity per second.
- The type is double.
- The unit is activity per sec.

## **Component: Sybase\_Engine\_Detail**

It contains information about the engines.

### **Dimensions**

#### Sybase Engine Detail Host Name

- The host name where the Sybase server resides.
- The type is string.

#### Engine Name (Unicode)

- The name of the engine.
- The type is string.

#### Sybase Engine Detail Row Number

- The row number in sample.
- The type is int.

#### Engine Name

- The name of the engine of the detailed statistics.
- The type is string. This is a key dimension.

#### Sybase Engine Detail Current Interval (Sec.)

- The number of seconds in the dbcc monitor sampling period.
- The type is int.

#### Sybase Engine Detail Server

- The name of the Sybase server.
- The type is string.

#### Sybase Engine Detail Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### Engine Number

- The engine number.
- The type is string.

#### Sybase Engine Detail Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

#### Sybase Engine Detail Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

### Metrics

#### TDS Packets Received Percent

- The percentage of server network packets received by this engine during the period.
- The type is double.
- The unit is percent.

#### TDS Bytes Sent (per Trans.)

- The number of network bytes sent per transaction by the engine during the period.
- The type is int.
- The unit is bytes.

#### Sybase Engine Detail Task Switch Percent

- The percentage of task context switches that occurred on a specific engine.
- The type is double.
- The unit is percent.

#### Sybase Engine Detail CPU Busy Percent

- The percentage of CPU busy during the sampling period.
- The type is double.
- The unit is percent.

#### TDS Packets Sent (per Trans.)

- The number of network packets sent per transaction by the engine during the period.
- The type is double.
- The unit is count.

#### TDS Packets Received (per Trans.)

- The number of network packets received per transaction by the engine during the period.

- The type is double.
- The unit is count.

#### CPU Yields Percent

- The percentage of server engine yields represented by the specified engine.
- The type is double.
- The unit is percent.

#### Sybase Engine Detail Total Transactions

- The number of transactions completed during the dbcc monitor sampling period.
- The type is int.
- The unit is transactions.

#### Engine Status

- The engine status.
- The type is string.
- The unit is status.

#### TDS Bytes Received (per Trans.)

- The number of network bytes received per transaction by the engine during the period.
- The type is int.
- The unit is bytes.

#### Engine Connections

- The count of the engine connections.
- The type is int.
- The unit is count.

#### TDS Bytes Received Percent

- The percentage of server network packets received by this engine during the period.
- The type is double.
- The unit is percent.

#### TDS Bytes Sent Percent

- The percentage of server network bytes sent by the engine during the period.
- The type is double.
- The unit is percent.

#### CPU Yields (per Trans.)

- The count number of occurrences that the CPU was released per transaction by an executing thread.
- The type is int.
- The unit is count.

#### Sybase Engine Detail Task Switch (per Trans.)

- The number of occurrences per transaction that the CPU switched between two threads.
- The type is double.

- The unit is count.

#### TDS Bytes Sent (per Sec.)

- The number of network bytes sent per second by the engine during the period.
- The type is int.
- The unit is bytes.

#### Sybase Engine Detail Maximum Outstanding IO Count

- The number of accesses to disk pending processing during the sampling period.
- The type is int.
- The unit is count.

#### Sybase Engine Detail TDS Bytes Received Count

- The number of network bytes received by the engine during the period.
- The type is int.
- The unit is count.

#### Sybase Engine Detail CPU Available Percent

- The percentage of CPU availability during the sampling period.
- The type is double.
- The unit is percent.

#### Sybase Engine Detail TDS Bytes Sent Count

- The number of network bytes sent by the engine during the period.
- The type is int.
- The unit is bytes.

#### Sybase Engine Detail Task Switch Count

- The number of occurrences that the CPU switched between two threads.
- The type is int.
- The unit is count.

#### TDS Bytes Received (per Sec.)

- The number of network bytes received per second by the engine during the period.
- The type is int.
- The unit is bytes.

#### TDS Packets Sent Percent

- The percentage of server network packets sent by the engine during the period.
- The type is double.
- The unit is percent.

#### Sybase Engine Detail Completed Disk IO Count

- The number of accesses to disk completed during the sampling period.
- The type is int.
- The unit is count.

#### TDS Packets Sent (per Sec.)

- The number of network packets sent per second by the engine during the period.
- The type is double.
- The unit is count.

#### Sybase Engine Detail Task Switch (per Sec.)

- The number of occurrences per second that the CPU switched between two threads.
- The type is int.
- The unit is count.

#### Sybase Engine Detail TDS Packets Received Count

- The number of network packets received by the engine during the period.
- The type is int.
- The unit is count.

#### Sybase Engine Detail TDS Packets Sent Count

- The number of network packets sent by the engine during the period.
- The type is int.
- The unit is count.

#### CPU Yields (per Sec.)

- The count number of occurrences that the CPU was released per second by an executing thread.
- The type is int.
- The unit is count.

#### Engine Idle Time Secs

- The engine idle time in seconds since last cycle.
- The type is int.
- The unit is secs.

#### Sybase Engine Detail CPU Yields Count

- The count number of occurrences that the cpu was released by an executing thread.
- The type is int.
- The unit is count.

#### TDS Packets Received (per Sec.)

- The number of network packets received per second by the engine during the period.
- The type is double.
- The unit is count.

### **Component: sybaseCacheDetail**

It contains sybase cache detail.

#### **Dimensions**

##### Sybase Cache Detail Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

#### Cache Name (Unicode)

- The cache name of the SQL server.
- The type is string.

#### Sybase Cache Detail Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

#### Sybase Cache Detail Host Name

- The host name where the Sybase server resides.
- The type is string.

#### SybaseCacheDetail Row Number

- The row number in sample.
- The type is int.

#### SybaseCacheDetail Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### SybaseCacheDetail Server

- The name of the Sybase server.
- The type is string.

#### Cache Name

- The cache name of the SQL server.
- The type is string. This is a key dimension.

### **Metrics**

#### Sybase Cache Detail LRU Buffer Use (per Trans.)

- The number of times per transaction that a LRU buffer was used.
- The type is double.
- The unit is count per trans.

#### Sybase Cache Detail Large IO Pages Used Count

- The count number of occurrences during the interval that a large IO page was used.
- The type is int.
- The unit is count.

#### Sybase Cache Detail LRU Buffer Use Percent

- The percentage of LRU buffers used.
- The type is double.
- The unit is percent.

#### Sybase Cache Detail Cache Miss Percent

- The percentage of cache search that a page was not found in the cache.



- The type is double.
- The unit is percent.

#### Sybase Cache Detail LRU Buffer Use Rate

- The number of times per second that a LRU buffer was used.
- The type is double.
- The unit is count per sec.

#### Sybase Cache Detail Cache Miss Count

- The count number of occurrences during the interval that a page was not found in the cache.
- The type is int.
- The unit is count.

#### Sybase Cache Detail Large IO Pages Used Percent

- The percentage of large IO pages cached that the pages were used.
- The type is double.
- The unit is percent.

#### Sybase Cache Detail Large IO Pages Used Rate

- The count number of occurrences per second that a large IO page was used.
- The type is double.
- The unit is occurrences per sec.

#### Sybase Cache Detail MRU Buffer Use (per Trans.)

- The number of times per transaction that an MRU buffer was used.
- The type is double.
- The unit is count per trans.

#### Sybase Cache Detail Large IO Pages Used (per Trans.)

- The count number of occurrences per transaction that a large IO page was used.
- The type is double.
- The unit is occurrences per trans.

#### Sybase Cache Detail Large IO Denied Percent

- The percentage of large IO requests that had been denied.
- The type is double.
- The unit is percent.

#### Sybase Cache Detail Large IO Pages Cached Rate

- The count number of occurrences per second that a large IO page was cached.
- The type is double.
- The unit is occurrences per sec.

#### Sybase Cache Detail MRU Buffer Use Rate

- The number of times per second that an MRU buffer was used.
- The type is double.
- The unit is count per sec.

#### Sybase Cache Detail Cache Miss (per Trans.)

- The count number of occurrences per transaction that a page was not found in the cache.
- The type is double.
- The unit is count per trans.

#### Sybase Cache Detail LRU Buffer Use Count

- The number of times during the interval that a LRU buffer was used.
- The type is int.
- The unit is count.

#### Sybase Cache Detail Cache Search Rate

- The count number of occurrences per second that a cache was searched.
- The type is double.
- The unit is count per sec.

#### Sybase Cache Detail Cache Search Count

- The count number of occurrences that a cache was searched.
- The type is int.
- The unit is count.

#### Sybase Cache Detail Large IO Performed Rate

- The count number of occurrences per second that a large IO was performed.
- The type is double.
- The unit is occurrences per sec.

#### Sybase Cache Detail Cache Miss Rate

- The count number of occurrences per second that a page was not found in the cache.
- The type is double.
- The unit is count per sec.

#### Sybase Cache Detail Large IO Performed (per Trans.)

- The count number of occurrences per transaction that a large IO was performed.
- The type is double.
- The unit is occurrences per trans.

#### Sybase Cache Detail MRU Buffer Use Count

- The number of times during the interval that an MRU buffer was used.
- The type is int.
- The unit is count.

#### Sybase Cache Detail Large IO Denied Count

- The count number of occurrences during the interval that a large IO was denied.
- The type is int.
- The unit is count.

#### Sybase Cache Detail Large IO Pages Cached Count

- The count number of occurrences during the interval that a large IO page was cached.
- The type is int.
- The unit is count.

#### Sybase Cache Detail Large IO Performed Percent

- The percentage of large IO requests that had been performed.
- The type is double.
- The unit is percent.

#### Sybase Cache Detail Cache Hit Count

- The count number of occurrences during the interval that a page was found in the cache.
- The type is int.
- The unit is count.

#### Sybase Cache Detail Cache Hit Rate

- The count number of occurrences per second that a page was found in the cache.
- The type is double.
- The unit is count per sec.

#### Sybase Cache Detail Large IO Pages Cached (per Trans.)

- The count number of occurrences per transaction that a large IO page was cached.
- The type is double.
- The unit is occurrences per trans.

#### Sybase Cache Detail MRU Buffer Use Percent

- The percentage of the MRU buffers used.
- The type is double.
- The unit is percent.

#### Sybase Cache Detail Large IO Denied (per Trans.)

- The count number of occurrences per transaction that a large IO was denied.
- The type is double.
- The unit is occurrences per trans.

#### Sybase Cache Detail Cache Search (per Trans.)

- The count number of occurrences per transaction that a cache was searched.
- The type is double.
- The unit is count per trans.

#### Sybase Cache Detail Cache Hit Percent

- The percentage of cache search that a page was found.
- The type is double.
- The unit is percent.

#### Sybase Cache Detail Large IO Denied Rate

- The count number of occurrences per second that a large IO was denied.
- The type is double.

- The unit is occurrences per sec.

Sybase Cache Detail Cache Hit (per Trans.)

- The count number of occurrences per transaction that a page was found in the cache.
- The type is double.
- The unit is count per trans.

Sybase Cache Detail Large IO Performed Count

- The count number of occurrences during the interval that a large IO was performed.
- The type is int.
- The unit is count.

### **Component: sybaseTaskDetail**

It contains sybase task detail.

#### **Dimensions**

Sybase Task Detail Host Name

- The host name where the Sybase server resides.
- The type is string.

SybaseTaskDetail Row Number

- The row number in sample.
- The type is int.

SybaseTaskDetail Server

- The name of the Sybase server.
- The type is string.

Sybase Task Detail Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

Sybase Task Detail Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

Task Switch Cause

- The name of the task switch cause (cache search misses, logical lock contention).
- The type is string. This is a key dimension.

SybaseTaskDetail Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### **Metrics**

Task Switch (per Trans.)

- The number of task switches per transaction.
- The type is double.

- The unit is count per trans.

Task Switch (per Sec.)

- The number of task switches per second.
- The type is double.
- The unit is count per sec.

Sybase Task Detail Task Switch Count

- The number of task switches during the sampling period.
- The type is int.
- The unit is count.

Task Switch Percent

- The percentage of task switches.
- The type is double.
- The unit is percent.

### **Component: sybaseProcessSummary**

It contains sybase process summary.

#### **Dimensions**

Sybase Process Summary Host Name

- The host name where SQLserver resides.
- The type is string.

SybaseProcessSummary Description

- The attribute name.
- The type is string.

SybaseProcessSummary Row Number

- The row number in sample.
- The type is int.

SybaseProcessSummary Server

- The name of the SQL Server.
- The type is string.

SybaseProcessSummary Value

- The attribute value.
- The type is string.

Sybase Process Summary Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

SybaseProcessSummary Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

## Metrics

### Server CPU Percent Application

- The percentage of CPU usage by Server Application processes, which exist at the sampling time. Note, the percentage value can be under reported if a process consumes high CPU, and it ends before the sample was taken. The value includes Sybase internal processes (to be comparable with Sybase server CPU), and it excludes the OS CPU reported in KOYPRCD CPUPCT.
- The type is double.
- The unit is percent.

### Percent Processes Bad

- The percentage of processes with a status of bad.
- The type is double.
- The unit is percent.

### Server CPU Pct System

- The percentage of CPU usage by Server System processes (CPUPCT - APPCPUPCT), it can be over reported because APPCPUPCT is under reported.
- The type is double.
- The unit is percent.

### Total Processes Stopped

- The number of processes with a status of stopped.
- The type is int.
- The unit is processes.

### Percent Processes Stopped

- The percentage of processes with a status of stopped.
- The type is double.
- The unit is percent.

### Sybase Process Summary Current Interval (Sec.)

- The number of seconds that have elapsed between the previous sample and the current sample.
- The type is int.
- The unit is secs.

### Percent Processes Infected

- The percentage of processes with a status infected.
- The type is double.
- The unit is percent.

### Total Processes Other Sleep

- The number of processes with a status of sleep for a reason, with exception for lock.
- The type is int.
- The unit is processes.

### Total Log Suspend

- The number of processes in log suspend.
- The type is int.
- The unit is processes.

#### Total Processes Infected

- The number of processes with a status of infected.
- The type is int.
- The unit is processes.

#### Total Processes Blocked

- The number of processes with a status of blocked.
- The type is int.
- The unit is processes.

#### Percent Processes Blocked

- The percentage of processes with a status of blocked.
- The type is double.
- The unit is percent.

#### Percent Processes Other Sleep

- The percentage of processes with a status of sleep for a reason, with exception for lock.
- The type is double.
- The unit is percent.

#### Percent Processes Sleeping

- The percentage of processes with a status of sleeping.
- The type is double.
- The unit is percent.

#### Total Processes

- The number of processes.
- The type is int.
- The unit is processes.

#### Total Processes Bad

- The number of processes with a status of bad.
- The type is int.
- The unit is processes.

#### Percent Processes Lock Sleep

- The percentage of processes with a status of lock sleep.
- The type is double.
- The unit is percent.

#### Total Processes Lock Sleep

- The number of processes with a status of lock sleep.
- The type is int.

- The unit is processes.

Sybase Process Summary Sample Timestamp

- The time when these data were collected.
- The type is timestamp.
- The unit is timestamp.

### **Component: sybaseJobSummary**

It contains sybase job summary.

#### **Dimensions**

Sybase Job Summary Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

Sybase Job Summary Host Name

- The host name where the SQL server resides.
- The type is string.

SybaseJobSummary Row Number

- The row number in sample.
- The type is int.

Sybase Job Summary Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

SybaseJobSummary Server

- The name of the SQL server.
- The type is string.

SybaseJobSummary Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### **Metrics**

Total Jobs

- The total jobs in the history.
- The type is int.
- The unit is jobs.

Succeeded Jobs Count

- The total jobs succeeded in the history.
- The type is int.
- The unit is jobs.

### **Component: sybaseLogDetail**

It contains sybase log detail.



## Dimensions

### Log Activity Name

- Indicates the name of the logging activity (Cache flushes, locks granted, locks waited).
- The type is string. This is a key dimension.

### Sybase Log Detail Host Name

- The host name where the Sybase server resides.
- The type is string.

### Sybase Log Detail Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

### SybaseLogDetail Server

- The ID of the Sybase server.
- The type is string.

### Sybase Log Detail Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

### SybaseLogDetail Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

### SybaseLogDetail Row Number

- The row number in sample.
- The type is int.

## Metrics

### Log Activity (per Sec.)

- The count number of occurrences of the activity per second.
- The type is double.
- The unit is activity per sec.

### Log Activity Count

- The count number of occurrences of the log activity during the sampling period.
- The type is int.
- The unit is count.

### Log Activity (per Trans.)

- The count number of occurrences of the activity per transaction.
- The type is double.
- The unit is activity per trans.

### Log Activity Percent

- The percentage of log cache activity represented by the specified activity.

- The type is double.
- The unit is percent.

**Component: sybasePhysicalDeviceDetail**

It contains sybase physical device detail.

**Dimensions**

Physical Name

- The name of the physical device used by the operating system.
- The type is string.

Device Name

- The name of the device recorded by the Sybase server.
- The type is string. This is a key dimension.

Physical Name (Unicode)

- The name of the physical device used by the operating system.
- The type is string.

SybasePhysicalDeviceDetail Server

- The name of the Sybase server.
- The type is string.

Sybase Physical Device Detail Host Name

- The host name where the Sybase server resides.
- The type is string.

SybasePhysicalDeviceDetail Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

Sybase Physical Device Detail Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

Sybase Physical Device Detail Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

Device Name (Unicode)

- The name of the device recorded by the Sybase server.
- The type is string. This is a key dimension.

SybasePhysicalDeviceDetail Row Number

- The row number in sample.
- The type is int.

## Metrics

### Device Lock Contention (per Sec.)

- The number of device locks that were denied per second.
- The type is double.
- The unit is locks per sec.

### Device Reads Count

- The count number of occurrences that the device was read from during the sampling period.
- The type is int.
- The unit is count.

### Device Writes Count

- The count number of occurrences that the device was written to during the sampling period.
- The type is int.
- The unit is writes.

### Device Lock Contention (per Trans.)

- The number of device locks that were denied per transaction.
- The type is double.
- The unit is locks per trans.

### Device Reads (per Sec.)

- The count number of occurrences that the device was read per second.
- The type is double.
- The unit is reads per sec.

### Device Total (per Trans.)

- The count number of occurrences of the device reads and the device writes per transaction.
- The type is double.
- The unit is count per trans.

### Device Reads (per Trans.)

- The count number of occurrences that the device was read per transaction.
- The type is double.
- The unit is reads per trans.

### Device Total Count

- The count number of occurrences of the device reads and the device writes during the sampling period.
- The type is int.
- The unit is count.

### Device Lock Contention Count

- The number of device locks that were denied during the sampling period.
- The type is int.
- The unit is locks.

#### Device Total Percent

- The percentage of server device activity which were directed to the device.
- The type is double.
- The unit is percent.

#### Device Lock Granted (per Trans.)

- The number of device locks that were granted per transaction.
- The type is double.
- The unit is locks per trans.

#### Device Lock Granted Percent

- The percentage of lock activity that was granted.
- The type is double.
- The unit is percent.

#### Device Writes (per Sec.)

- The count number of occurrences that the device was written per second.
- The type is double.
- The unit is writes per sec.

#### Device Lock Contention Percent

- The percentage of lock activity that was denied.
- The type is double.
- The unit is percent.

#### Device Lock Granted Count

- The number of device locks that were granted during the sampling period.
- The type is int.
- The unit is count.

#### Device Writes Percent

- The percentage of device writes activity.
- The type is double.
- The unit is percent.

#### Device Writes (per Trans.)

- The count number of occurrences that the device was written per transaction.
- The type is double.
- The unit is writes per trans.

#### Device Lock Granted (per Sec.)

- The number of device locks that were granted per second.
- The type is double.
- The unit is locks per sec.

#### Device Total (per Sec.)

- The count number of occurrences of the device reads and the device writes per second.

- The type is double.
- The unit is count per sec.

#### Device Reads Percent

- The percentage of the device reads activity.
- The type is double.
- The unit is percent.

### **Component: sybaseEngineSummary**

It contains sybase engine summary.

#### **Dimensions**

##### SybaseEngineSummary Server

- The name of the Sybase server.
- The type is string.

##### Sybase Engine Summary Host Name

- The host name where the Sybase server resides.
- The type is string.

##### SybaseEngineSummary Row number

- The row number in sample.
- The type is int.

##### SybaseEngineSummary Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

##### Sybase Engine Summary Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

##### Engine Count

- The number of engines assigned to the server.
- The type is string.

##### Sybase Engine SummarySample Timestamp

- The time when these data were collected.
- The type is timestamp.

#### **Metrics**

##### Total Transactions

- The number of transactions that completed during the dbcc monitor sampling period.
- The type is int.
- The unit is transactions.

##### CPU Busy Percent

- The percentage of time the server was busy executing tasks during the sampling period.

- The type is double.
- The unit is percent.

#### CPU Yields Count

- The number of yields by the server engine to the operating system.
- The type is int.
- The unit is count.

#### Maximum Outstanding IO Count

- The number of accesses to disk pending processing during the sampling period.
- The type is int.
- The unit is count.

#### Sybase Engine Summary Current Interval (Sec.)

- The number of seconds in the dbcc monitor sampling period.
- The type is int.
- The unit is secs.

#### Task Switch Count

- The number of changes by the server from one user task to another during the sampling period.
- The type is int.
- The unit is count.

#### Completed Disk IO Count

- The number of accesses to disk completed during the sampling period.
- The type is int.
- The unit is count.

#### TDS Bytes Received Count

- The number of bytes the engine received during the sampling period.
- The type is int.
- The unit is bytes.

#### TDS Packets Received Count

- The number of packets the engine received during the sampling period.
- The type is int.
- The unit is packets.

#### TDS Packets Sent Count

- The number of network packets sent by the engine during the sampling period.
- The type is int.
- The unit is packets.

#### CPU Available Percent

- The percentage of time the server was not busy executing tasks during the sampling period.
- The type is double.
- The unit is percent.

TDS Bytes Sent Count

- The number of network bytes sent by the engine during the sampling period.
- The type is int.
- The unit is bytes.

### **Component: SybaseServerSummary**

Sybase\_Server\_Summary.

#### **Dimensions**

Sybase Server Summary Server Type

- The type of Sybase Server.
- The type is string.

SybaseServerSummary Server

- The name of the Sybase server.
- The type is string.

Sybase Server Summary Collection Status

- Indicates the status of Data Collection on remote node (Active, Inactive, DB\_Connect\_Fail, DC\_Connect\_Fail).
- The type is string.

Sybase Server Summary Data Cache Size (KB)

- The number of kilobytes (KB) allocated for the data buffer cache.
- The type is int.

Sybase Server Summary Procedure Cache Size (KB)

- The number of kilobytes (KB) allocated for the procedure cache.
- The type is int.

Sybase Server Summary Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

SybaseServerSummary Description

- The attribute name.
- The type is string.

Sybase Server Summary Sample Timestamp

- The timestamp that indicates the date and time the product collected the sample data for the Sybase server.
- The type is timestamp.

SybaseServerSummary Value

- The attribute value.
- The type is string.

Sybase Server Summary Host Name

- The host name where the Sybase server resides.
- The type is string.

#### SybaseServerSummary Row Number

- The row number in sample.
- The type is int.

#### Sybase Server Summary Server Version

- The version of the Sybase server.
- The type is string.

#### SybaseServerSummary Originnode

- The key for the table in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

### Metrics

#### Sybase Server Summary Time Since Startup (Min.)

- The number of minutes that have elapsed since the server was started.
- The type is int.
- The unit is minutes.

#### Sybase Server Summary Server CPU Percent

- The percentage of CPU usage by the Sybase server process measured on the host.
- The type is double.
- The unit is percent.

#### Sybase Server Summary Total OS CPU Percent

- The percentage of CPU usage by all the processes on the host.
- The type is double.
- The unit is percent.

#### Sybase Server Summary Server Status

- Indicates the status of the SQL server (Active, Inactive or Unknown).
- The type is string.
- The unit is status.

#### Pct Max Locks

- The percentage of locks on resources of the maximum number of locks allowed by the Sybase server.
- The type is double.
- The unit is percent.

#### Sybase Server Summary Current Interval (Sec.)

- The number of seconds that have elapsed between the previous sample and the current sample.
- The type is int.
- The unit is secs.



## **Component: sybaseLockSummary**

It contains sybase lock summary.

### **Dimensions**

Sybase Lock Summary Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

Sybase Lock Summary Host Name

- The host name where the Sybase server resides.
- The type is string.

SybaseLockSummary Row Number

- The row number in sample.
- The type is int.

Sybase Lock Summary Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

SybaseLockSummary Server

- The name of the Sybase server.
- The type is string.

SybaseLockSummary Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

### **Metrics**

Lock Contention (per Trans.)

- The average number of requests per transaction waiting for locks on resources to be released during the sampling period.
- The type is double.
- The unit is requests per trans.

Lock Promotes (per Sec.)

- The number of lock promotions per second.
- The type is double.
- The unit is lock promotions per sec.

Lock Promotes (per Trans.)

- The number of lock promotions per transaction.
- The type is double.
- The unit is lock promotions per trans.

Lock Request (per Sec.)

- The number of lock calls per second.

- The type is double.
- The unit is lock calls per sec.

#### Lock Promotes Count

- The number of lock promotions during the sampling period.
- The type is int.
- The unit is count.

#### Deadlocks (per Sec.)

- The number of deadlocks per second.
- The type is double.
- The unit is deadlocks per sec.

#### Lock Promotes Exclusive Percent

- The percentage of exclusive lock promotions.
- The type is double.
- The unit is percent.

#### Deadlocks Count

- The number of deadlocks during the sampling period.
- The type is int.
- The unit is count.

#### Lock Promotes Exclusive (per Trans.)

- The number of exclusive lock promotions per transaction.
- The type is double.
- The unit is lock promotions per trans.

#### Lock Request Count

- The number of lock calls during the sampling period.
- The type is int.
- The unit is count.

#### Deadlocks (per Trans.)

- The number of deadlocks per transaction.
- The type is double.
- The unit is deadlocks per trans.

#### Lock Promotes Shared Percent

- The percentage of shared lock promotions.
- The type is double.
- The unit is percent.

#### Deadlocks Percent

- The percentage of lock requests that resulted in deadlocks for the server during the sampling period.
- The type is double.

- The unit is percent.

#### Lock Request (per Trans.)

- The number of lock calls per transaction.
- The type is double.
- The unit is lock calls per trans.

#### Lock Wait Time

- The longest waiting time in seconds for a lock during the sample period.
- The type is int.
- The unit is secs.

#### Lock Promotes Exclusive (per Sec.)

- The number of exclusive lock promotions per second.
- The type is double.
- The unit is lock promotions per sec.

#### Database Name (Unicode)

- The name of the database that is waiting on a lock for the longest period of time.
- The type is string.
- The unit is Name.

#### Lock Contention (per Sec.)

- The average number of lock requests per second waiting for locks on resources to be released during the sampling period.
- The type is double.
- The unit is requests per sec.

#### Lock Promotes Shared Count

- The number of shared lock promotions during the sampling period.
- The type is int.
- The unit is count.

#### ID of Longest Waiting Lock

- The ID of the longest waiting lock.
- The type is string.
- The unit is ID.

#### Lock Promotes Exclusive Count

- The number of exclusive lock promotions during the sampling period.
- The type is int.
- The unit is count.

#### Lock Promotes Shared (per Sec.)

- The number of shared lock promotions per second.
- The type is double.
- The unit is lock promotions per sec.

#### Lock Contention Count

- The number of requests waiting for locks on resources to be released during the sampling period.
- The type is int.
- The unit is count.

#### Lock Contention Percent

- The percentage of requests waiting for locks on resources to be released during the sampling period.
- The type is double.
- The unit is percent.

#### Lock Promotes Shared (per Trans.)

- The number of shared lock promotions per transaction.
- The type is double.
- The unit is lock promotions per trans.

#### Available Locks

- Total number of locks available.
- The type is int.
- The unit is locks.

### **Component: sybaseDatabaseSummary**

SybaseDatabaseSummary.

#### **Dimensions**

##### SybaseDatabaseSummary Description

- The attribute name.
- The type is string.

##### SybaseDatabaseSummary Row Number

- The row number in sample.
- The type is int.

##### SybaseDatabaseSummary Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

##### SybaseDatabaseSummary Value

- The attribute value.
- The type is string.

##### SybaseDatabaseSummary Server

- The name of the Sybase server.
- The type is string.

##### Sybase Database Summary Host Name

- The host name where the Sybase server resides.

- The type is string.

#### Sybase Database Summary Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

#### Sybase Database Summary Sample Timestamp

- The timestamp that indicates the date and time the product collected the sample for the server.
- The type is timestamp.

### **Metrics**

#### Current Interval (Sec.)

- The number of seconds that have elapsed between the previous sample and the current sample.
- The type is int.
- The unit is sec.

#### Total DBs Single User

- The number of databases with a status of Single User mode.
- The type is int.
- The unit is dbs.

#### Total DBs DBO Only

- The number of databases with a status of DBO only.
- The type is int.
- The unit is dbs.

#### Minimum Pct Log Freespace

- The minimum percentage of free space in the database logs.
- The type is double.
- The unit is percent.

#### Minimum Pct Data Freespace

- The minimum percentage of free space in the database.
- The type is double.
- The unit is percent.

#### Total Databases

- The total number of databases in the server.
- The type is int.
- The unit is databases.

#### Total DBs Read Only

- The number of databases with a status of read only status.
- The type is int.
- The unit is dbs.

#### Total DBs in Error

- The number of databases with an error status.
- The type is int.
- The unit is databases.

#### Total DBs No Free Space Accounting

- The number of databases with a status of No Free Space Accounting.
- The type is int.
- The unit is dbs.

#### Number of Databases Open

- The number of databases open.
- The type is int.
- The unit is dbs.

### **Component: sybaseProcessDetail**

It contains sybase process detail.

#### **Dimensions**

##### Program Name

- The program name.
- The type is string.

##### Blocking Process ID (Unicode)

- The process ID of the process which is blocking this process.
- The type is string.

##### Server User Name (Unicode)

- The server user ID.
- The type is string.

##### Sybase Process Detail

- The name of the executing command.
- The type is string.

##### Login Name

- The process name.
- The type is string.

##### Server User Name

- The server user ID.
- The type is string.

##### Program Name (Unicode)

- The program name.
- The type is string.

##### Client Process ID

- The process ID in the client host.

- The type is string. This is a key dimension.

#### Client Group ID

- The group ID of the user.
- The type is string.

#### Client Host Name

- The host name of the client where the command was issued.
- The type is string. This is a key dimension.

#### Client User ID

- The user ID who executed the command.
- The type is string.

#### Command (Unicode)

- The name of the executing command.
- The type is string.

#### Transaction Name (Unicode)

- The transaction name.
- The type is string.

#### Transaction Name

- The transaction name.
- The type is string.

#### Sybase Process Detail Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

#### SybaseProcessDetail Server

- The name of the SQL Server.
- The type is string.

#### Sybase Process Detail Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

#### Sybase Process Detail Host Name

- The host name where SQL server resides.
- The type is string.

#### Process ID

- The process ID.
- The type is int. This is a key dimension.

#### Client User ID (Unicode)

- The user ID who executed the command.

- The type is string.

#### SybaseProcessDetail Row Number

- The row number in sample.
- The type is int.

#### Sybase Process Detail Database Name (Unicode)

- The database name.
- The type is string.

#### OS Process ID

- The OS process ID.
- The type is string.

#### Sybase Process Detail Blocking Process ID

- The process ID of the process that is blocking this process.
- The type is string.

#### SybaseProcessDetail Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### Sybase Process Detail Database Name

- The database name.
- The type is string.

#### Login Name (Unicode)

- The process name.
- The type is string.

#### Client Group ID (Unicode)

- The group ID of the User.
- The type is string.

### **Metrics**

#### Total Memory Alloc

- The total amount of memory allocated in kilobytes.
- The type is int.
- The unit is kb.

#### Current CPU Pct Used

- The percentage of CPU usage since the server was started.
- The type is double.
- The unit is percent.

#### Total Disk IO

- The number of disk operations on reads and writes.



- The type is int.
- The unit is count.

#### Network Packet Size (Bytes)

- The network package size in bytes.
- The type is int.
- The unit is bytes.

#### Sybase Process Detail Time Blocked (Sec.)

- The total amount of time in seconds that the process has been blocked.
- The type is double.
- The unit is sec.

#### Total CPU Time (Sec.)

- The CPU time usage in seconds, both external CPU charged by the OS to the process, and internal CPU executed by the Sybase server for the processes.
- The type is double.
- The unit is sec.

#### Process Status

- Indicates the process status (lock sleep, infected, bad, stopped or log suspend).
- The type is string.
- The unit is status.

### **Component: sybaseConfiguration**

It contains sybase configuration.

#### **Dimensions**

##### Minimum Value

- The minimum value permitted.
- The type is string.

##### Sybase Configuration Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

##### SybaseConfiguration Row Number

- The row number in sample.
- The type is int.

##### Config Parameter

- The configuration parameter name.
- The type is string. This is a key dimension.

##### Sybase Configuration Host Name

- The host name where the Sybase server resides.
- The type is string.

#### Run Value

- The value of the parameter.
- The type is string.

#### Maximum Value

- The maximum value permitted.
- The type is string.

#### Config Parameter (Unicode)

- The configuration parameter name.
- The type is string.

#### Config Value (Unicode)

- The configured value of the parameter.
- The type is string.

#### Run Value (Unicode)

- The parameter value.
- The type is string.

#### SybaseConfiguration Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### Sybase Configuration Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

#### Config Value

- The configured value of the parameter.
- The type is string.

#### SybaseConfiguration Server

- The name of the Sybase server.
- The type is string.

#### Parm Type

- Indicates the parameter type. Parameter type of Dynamic(1) indicates that the server restart is not required. Parameter type of Static(0) indicates that the server restart is required to take effect.
- The type is string.

### **Component: Sybase\_Cache\_Summary**

It contains information about Sybase Cache Summary.

#### **Dimensions**

##### Sybase Cache Summary Server

- The name of the Sybase server.

- The type is string.

#### Sybase Cache Summary Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

#### Sybase Cache Summary Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

#### Sybase Cache Summary Row Number

- The row number in sample.
- The type is int.

#### Sybase Cache Summary Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### Sybase Cache Summary Host Name

- The host name where the Sybase server resides.
- The type is string.

### **Metrics**

#### Large IO Denied Count

- The count number of occurrences during the interval that a large IO was denied.
- The type is int.
- The unit is count.

#### Large IO Performed (per Trans.)

- The count number of occurrences per transaction that a large IO was performed.
- The type is double.
- The unit is count.

#### Cache Search (per Trans.)

- The count number of occurrences per transaction that the cache was searched.
- The type is double.
- The unit is count.

#### Large IO Pages Cached (per Trans.)

- The count number of occurrences per transaction that a large IO page was cached.
- The type is double.
- The unit is count.

#### Cache Hit (per Trans.)

- The count of number occurrences per transaction that a page was found in the cache.
- The type is double.
- The unit is count.

#### LRU Buffer Use Percent

- The percentage of LRU buffer used.
- The type is double.
- The unit is percent.

#### Large IO Pages Used Percent

- The percentage of large IO cached pages were used.
- The type is double.
- The unit is percent.

#### LRU Buffer Use Count

- The number of times during the interval that a LRU buffer was used.
- The type is int.
- The unit is count.

#### Large IO Denied (per Trans.)

- The count number of occurrences per transaction that a large IO was denied.
- The type is double.
- The unit is count.

#### Large IO Pages Cached Rate

- The count number of occurrences per second that a large IO page was cached.
- The type is double.
- The unit is count.

#### Cache Search Count

- The count number of occurrences that the cache was searched.
- The type is int.
- The unit is count.

#### MRU Buffer Use Count

- The number of times during the interval that a MRU buffer was used.
- The type is int.
- The unit is count.

#### LRU Buffer Use Rate

- The number of times per second that a LRU buffer was used.
- The type is double.
- The unit is rate.

#### Large IO Performed Count

- The count number of occurrences during the interval that a large IO was performed.
- The type is int.
- The unit is count.

#### Cache Hit Rate

- The count number of occurrences per second that a page was found in the cache.

- The type is double.
- The unit is count.

#### Cache Miss (per Trans.)

- The count number of occurrences per transaction that a page was not found in the cache.
- The type is double.
- The unit is count.

#### MRU Buffer Use Rate

- The number of times per second that a MRU buffer was used.
- The type is double.
- The unit is rate.

#### Large IO Performed Percent

- The percentage of large IO requests performed by the server during the current interval.
- The type is double.
- The unit is percent.

#### Large IO Pages Used Count

- The count number of occurrences during the interval that a large IO page was used.
- The type is int.
- The unit is count.

#### Cache Hit Percent

- The percentage of cache search that a page was found in the cache.
- The type is double.
- The unit is percent.

#### Large IO Pages Cached Count

- The count number of occurrences during the interval that a large IO page was cached.
- The type is int.
- The unit is count.

#### Large IO Denied Rate

- The count number of occurrences per second that a large IO was denied.
- The type is double.
- The unit is count.

#### Cache Search Rate

- The count number of occurrences per second that the cache was searched.
- The type is double.
- The unit is count.

#### Cache Miss Rate

- The count number of occurrences per second that a page was not found in the cache.
- The type is double.
- The unit is count.

#### Large IO Denied Percent

- The percentage of large IO requests that had been denied.
- The type is double.
- The unit is percent.

#### Large IO Pages Used (per Trans.)

- The count number of occurrences per transaction that a large IO page was used.
- The type is double.
- The unit is count.

#### Large IO Pages Used Rate

- The count number of occurrences per second that a large IO page was used.
- The type is double.
- The unit is count.

#### Cache Miss Percent

- The percentage of cache search that a page was not found in the cache.
- The type is double.
- The unit is percent.

#### Cache Hit Count

- The count of number occurrences during the interval that a page was found in the cache.
- The type is int.
- The unit is count.

#### Cache Miss Count

- The count number of occurrences during the interval that a page was not found in the cache.
- The type is int.
- The unit is count.

#### LRU Buffer Use (per Trans.)

- The number of times per transaction that a LRU buffer was used.
- The type is double.
- The unit is rate.

#### MRU Buffer Use (per Trans.)

- The number of times per transaction that a MRU buffer was used.
- The type is double.
- The unit is rate.

#### MRU Buffer Use Percent

- The percentage of MRU buffer used.
- The type is double.
- The unit is percent.

#### Large IO Performed Rate

- The count number of occurrences per second that a large IO was performed.
- The type is double.
- The unit is count.

### **Component: Sybase\_Job\_Detail**

It contains information about Sybase Job Detail.

#### **Dimensions**

Sybase Job Detail Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

Job Id

- The SQL server scheduled job ID.
- The type is int. This is a key dimension.

Sybase JobDetail Row Number

- The row number in sample.
- The type is int.

Sybase Job Detail Host Name

- The host name where the SQL server resides.
- The type is string.

Sybase Job Detail Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

Sybase Job Detail Server

- The name of the SQL server.
- The type is string. This is a key dimension.

Job Name (Unicode)

- The Sybase server job name.
- The type is string. This is a key dimension.

Job Description

- The description of the job.
- The type is string.

Sybase Job Detail Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

Job Owner

- The name of the job owner.
- The type is string.

## Metrics

### Schedule Interval

- The schedule interval between the scheduled job repeats.
- The type is string.
- The unit is ms.

### Last Run Outcome

- Indicates the previous job execution status (Unknown, Success, Failure, Ext\_Failure, Normal).
- The type is string.
- The unit is status.

### Schedule Name (Unicode)

- The Sybase server schedule name.
- The type is string.
- The unit is name.

### Job Enabled

- Indicates whether the job is enabled to run.
- The type is int.
- The unit is count.

### Job State

- The Sybase server job state (waiting, queued, busy, runnable, running, ncompleting, completed, terminating, terminated, timing\_out, timed\_out, missed).
- The type is string.
- The unit is status.

### Last Run Timestamp

- The timestamp of the last job execution.
- The type is timestamp.
- The unit is timestamp.

## Component: sybaseLockConflictDetail

It contains Sybase\_Lock\_Conflict\_Detail.

### Dimensions

#### SybaseLockConflictDetail Row Number

- The row number in sample.
- The type is int.

#### Server User ID (Unicode)

- The server user ID.
- The type is string.

#### Database Id

- The ID of the database.
- The type is int.



#### SybaseLockConflictDetail Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

#### Sybase Lock Conflict Detail Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

#### SybaseLockConflictDetail Server

- The name of the SQL server.
- The type is string.

#### Sybase Lock Conflict Detail Sample Timestamp

- The timestamp that indicates the date and time the product collected the data.
- The type is timestamp.

#### Sybase Lock Conflict Detail Database Name

- The database name.
- The type is string. This is a key dimension.

#### Sybase Lock Conflict Detail Host Name

- The host name where the server resides.
- The type is string.

#### Sybase Lock Conflict Detail Database Name (Unicode)

- The database name.
- The type is string.

### **Metrics**

#### Blocking Process ID

- The process which is blocking the lock request.
- The type is int.
- The unit is ID.

#### Time Blocked (Sec.)

- The total time (in seconds) the process has been blocked.
- The type is double.
- The unit is sec.

#### Server User ID

- The server user ID.
- The type is string.
- The unit is ID.

#### Requestor Process ID

- The process which is requesting for the lock.
- The type is int.

- The unit is ID.

### **Component: sybaseLockDetail**

It contains sybase lock detail.

#### **Dimensions**

Sybase Lock Detail Host Name

- The host name where the Sybase server resides.
- The type is string.

SybaseLockDetail Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

SybaseLockDetail Server

- The name of the Sybase server.
- The type is string.

Sybase Lock Detail Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

Lock Type

- The lock types (ex-table, sh-page, address).
- The type is string. This is a key dimension.

SybaseLockDetail Row Number

- The row number in sample.
- The type is int.

Sybase Lock Detail Sample Timestamp

- The time when these data were collected.nOPTION: ALIGNCENTER.
- The type is timestamp.

#### **Metrics**

Lock Deadlocked (per Sec.)

- The total number of deadlocked lock requests per second.
- The type is double.
- The unit is lock requests per sec.

Lock Granted (per Trans.)

- The total number of granted lock requests per transaction.
- The type is double.
- The unit is requests per trans.

Sybase Lock Detail Lock Contention (per Trans.)

- The total number of denied lock requests per transaction.
- The type is double.

- The unit is requests per trans.

#### Lock Granted Percent

- The percentage of lock requests granted during the current interval.
- The type is double.
- The unit is percent.

#### Lock Total Count

- The total number of lock requests during the sampling period.
- The type is int.
- The unit is count.

#### Lock Granted Count

- The total number of granted lock requests during the sampling period.
- The type is int.
- The unit is count.

#### Sybase Lock Detail Lock Contention Percent

- The percentage of denied lock requests.
- The type is double.
- The unit is percent.

#### Lock Granted (per Sec.)

- The total number of granted lock requests per second.
- The type is double.
- The unit is requests per sec.

#### Lock Deadlocked Percent

- The percentage of deadlocked lock requests.
- The type is double.
- The unit is percent.

#### Sybase Lock Detail Lock Contention Count

- The total number of denied lock requests during the sampling period.
- The type is int.
- The unit is count.

#### Lock Deadlocked Count

- The total number of deadlocked lock requests during the sampling period.
- The type is int.
- The unit is locks.

#### Sybase Lock Detail Lock Contention (per Sec.)

- The total number of denied lock requests per second.
- The type is double.
- The unit is requests per sec.

Lock Total Percent

- The percentage of lock requests represented by the specified lock type.
- The type is double.
- The unit is percent.

Lock Total (per Sec.)

- The total number of lock requests per second.
- The type is double.
- The unit is requests per sec.

Lock Deadlocked (per Trans.)

- The total number of deadlocked lock requests per transaction.
- The type is double.
- The unit is lock requests per trans.

Lock Total (per Trans.)

- The total number of lock requests per transaction.
- The type is double.
- The unit is requests per trans.

### **Component: sybaseStatisticsDetail**

It contains sybase statistics detail.

#### **Dimensions**

SybaseStatisticsDetail Server

- The name of the Sybase server.
- The type is string.

Sybase Statistics Detail Hub Timestamp

- The time when this data was inserted at hub.
- The type is timestamp.

SybaseStatisticsDetail Row Number

- The row number in sample.
- The type is int.

SybaseStatisticsDetail Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

Statistic Name

- The name of the statistic.
- The type is string. This is a key dimension.

Sybase Statistics Detail Host Name

- The host name where the Sybase server resides.
- The type is string.

#### Sybase Statistics Detail Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

### **Metrics**

#### Maximum Seen

- The maximum value per second since startup.
- The type is int.
- The unit is value.

#### Current Value

- The value during the current interval.
- The type is int.
- The unit is value.

#### Minimum Seen

- The minimum value per Second since startup.
- The type is int.
- The unit is value.

#### Average Value (per Sec.)

- The average value per second since startup.
- The type is double.
- The unit is value.

#### Total Since Startup

- The cumulative value since startup.
- The type is int.
- The unit is value.

### **Component: sybaseLogSummary**

It contains sybase log summary.

#### **Dimensions**

##### SybaseLogSummary Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

##### Sybase Log Summary Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

##### SybaseLogSummary Row Number

- The row number in sample.
- The type is int.

##### Sybase Log Summary Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

#### SybaseLogSummary Server

- The name of the Sybase server.
- The type is string.

#### Sybase Log Summary Host Name

- The host name where the Sybase server resides.
- The type is string.

### Metrics

#### PLC Semaphore Requests Count

- The number of Private Log Cache, or PLC, semaphore requests during the sampling period.
- The type is int.
- The unit is count.

#### Private Log Cache Flushes Count

- The number of private log cache flushes during the sampling period.
- The type is int.
- The unit is count.

#### Transaction Log Writes Count

- The number of transaction log writes to disk during the sampling period.
- The type is int.
- The unit is count.

#### Maximum Private Log Cache Size (KB)

- The maximum size in kilobytes of any private log cache.
- The type is int.
- The unit is kb.

#### Private Log Cache Log Records Count

- The number of pages for private log caches during the sampling period.
- The type is int.
- The unit is pages.

#### Log Semaphore Requests Count

- The number of requests waiting for locks on resources to be released during the sampling period.
- The type is int.
- The unit is count.

#### Transaction Log Allocates Count

- The number of page allocations to the transaction log during the sampling period.
- The type is int.
- The unit is pages.

**Component: sybaseText**

It contains sybase text.

**Dimensions**

SybaseText Row Number

- The row number in sample.
- The type is int.

SybaseText Server

- The name of the SQL Server.
- The type is string.

SybaseText Text

- The SQL text.
- The type is string.

Sequence Num

- The SQL sequence number.
- The type is int. This is a key dimension.

SybaseText Originnode

- The key for the table view in the format SYB.serverid.hostname.
- The type is string. This is a key dimension.

Sybase Text Process ID

- The Sybase process ID.
- The type is int. This is a key dimension.

Text (Unicode)

- The SQL text.
- The type is string.

Database Name

- The database name.
- The type is string.

Sybase Text Hub Timestamp

- The time when this data was inserted at the hub.
- The type is timestamp.

Sybase Text Host Name

- The host name where the server resides.
- The type is string.

Sybase Text Detail Sample Timestamp

- The time when these data were collected.
- The type is timestamp.

Database Name. (Unicode)

- The database name.
- The type is string.

## Windows OS agent metrics

The metrics for Windows OS agent resource types collect data for monitoring with IBM Cloud App Management. Every Windows OS agent resource type defines a set of dimensions and metrics. The descriptions provide such information as data type, dimension key, and metric unit.

### **windowsOS**

Windows Operating System.

#### **Dimensions**

Computer Name Unicode

- The commonly used product name (unicode). This attribute value is set by the computer system manufacturer.
- The type is string.

Computer System Description

- The computer system description. This attribute value is set by the computer system manufacturer.
- The type is string.

Computer Hostname

- The host name of the host computer. This value is resolved from the host IP address.
- The type is string. This is a key dimension.

Computer Vendor Unicode

- The supplier name of the product (unicode). Corresponds to the Vendor property in the product object in the DMTF Solution Exchange Standard. This attribute value is set by the computer system manufacturer.
- The type is string.

Computer Domain Name

- The fully qualified domain name of the host computer. This value is resolved from the host IP address.
- The type is string.

Computer ID Number

- This attribute value can be equal to the BIOS serial number for some Manufacturers. Typically this occurs when the computer system manufacturer and BIOS manufacturer are the same. This attribute value is set by the computer system manufacturer.
- The type is string.

Computer System Description Unicode

- The computer system description (unicode). This attribute value is set by the computer system manufacturer.
- The type is string.

System Name



- The managed system name. The format is *hostname : agent\_code* . Examples include *spark:KNT* or *deux.raleigh.ibm.com:KNT*.
- The type is string.

#### Computer Name

- The commonly used product name. This attribute value is set by the computer system manufacturer.
- The type is string.

#### Computer UUID Number

- The universally unique identifier (UUID) for this product. A UUID is a 128-bit identifier that is guaranteed to be different from other generated UUIDs. If a UUID is not available, a UUID of all zeros is used.
- The type is string.

#### Computer Version Unicode

- The product version information (unicode). Corresponds to the Version property in the product object in the DMTF Solution Exchange Standard. This attribute value is set by the computer system manufacturer.
- The type is string.

#### Computer Vendor

- The supplier name of the product. Corresponds to the Vendor property in the product object in the DMTF Solution Exchange Standard. This attribute value is set by the computer system manufacturer.
- The type is string.

#### Computer Version

- The product version information. Corresponds to the Version property in the product object in the DMTF Solution Exchange Standard. This attribute value is set by the computer system manufacturer.
- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

### **Component: Log File Profiles**

Log File Profiles.

#### **Dimensions**

Log File Profiles Subnode Affinity

- The affinity for the subnode agent.
- The type is string.

Log File Profiles Subnode Description

- User supplied description of this subnode, specified via the SubnodeDescription key in the config file.
- The type is string.

#### Log File Profiles Subnode Type

- The Node Type of this subnode.
- The type is string. This is a key dimension.

#### Log File Profiles Error Code

- The error code that is associated with the query.
- The type is int.

#### Log File Profiles Subnode Version

- The Version of the subnode agent.
- The type is string.

#### Log File Profiles Subnode MSN

- The Managed System Name of the subnode agent.
- The type is string. This is a key dimension.

#### Log File Profiles Subnode Resource Name

- The Resource Name of the subnode agent.
- The type is string.

#### Log File Profiles System Name

- This is the managed system name of the agent.
- The type is string.

#### Log File Profiles Object Status

- The status of the performance object.
- The type is int.

#### Log File Profiles Timestamp

- This is the local time when the data was collected.
- The type is timestamp.

#### Log File Profiles Subnode Config File

- The path to the discovered configuration file that caused this subnode to be created.
- The type is string.

### **Component: Active Tasks running on VCenter Server**

Use the MSMQ Sessions data set to monitor session statistics. MSMQ (Microsoft Message Queue) Sessions is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### MSMQ Session Session Name

- The IP address of the computer in session with MSMQ. For example, MBROWN2.
- The type is string. This is a key dimension.

##### MSMQ Session Info Store Parameter

- A special one for the front end to use as a column header.

- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### node

- The managed system name. The format should be *hostname : agent\_code* . Examples include spark:KNT or deux.raleigh.ibm.com:KNT.
- The type is string.

#### MSMQ Session Info Store Value

- A special one for the front end to use as a column header.
- The type is string.

### Metrics

#### MSMQ Session Incoming Bytes

- The total number of bytes that were received through the selected session.
- The type is int.
- The unit is bytes.

#### MSMQ Session Outgoing Bytes

- The total number of bytes that were sent through the selected session.
- The type is int.
- The unit is bytes.

#### MSMQ Session Outgoing Messages/sec

- The rate that MSMQ messages are leaving per second through the selected session.
- The type is int.
- The unit is messages/second.

#### MSMQ Session Incoming Messages

- The total number of messages that were received through the selected session.
- The type is int.
- The unit is messages.

#### MSMQ Session Replication Requests Received

- The total number of replication requests received.
- The type is int.
- The unit is requests.

#### MSMQ Session Sync Replies

- The total number of sync requests that were answered.
- The type is int.
- The unit is requests.

#### MSMQ Session Access to the Server

- The total number of times the MSMQ Information Store was accessed.
- The type is int.
- The unit is accesses.

#### MSMQ Session Sync Requests

- The total number of sync requests received.
- The type is int.
- The unit is requests.

#### MSMQ Session Write Requests Sent

- The total number of write requests sent.
- The type is int.
- The unit is requests.

#### MSMQ Session Incoming Bytes/sec

- The rate that MSMQ messages are entering through the selected session.
- The type is int.
- The unit is bytes/second.

#### MSMQ Session Outgoing Messages

- The total number of messages that were sent through the selected session.
- The type is int.
- The unit is messages.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### MSMQ Session Errors Returned to Application

- The total number of MSMQ Information Store accesses that resulted in error replies by the Information Store.
- The type is int.
- The unit is accesses.

#### MSMQ Session Replication Requests Sent

- The total number of replication requests sent.
- The type is int.
- The unit is requests.

#### MSMQ Session Incoming Messages/sec

- The rate that MSMQ messages are entering per second through the selected session.
- The type is int.
- The unit is messages/second.

#### MSMQ Session Outgoing Bytes/sec

- The rate that MSMQ messages are leaving per second through the selected session.
- The type is int.
- The unit is bytes/second.

### **Component: Active Tasks running on VCenter Server**

Use the Physical Disk data set to create situations that monitor information about fixed and hard disk drives. Physical disk is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

node

- The managed system name.
- The type is string.

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

Disk Number

- The number of a physical disk. Calculated as the left part of the Disk Name.
- The type is string.

Physical Disk Name

- The name of a physical disk.
- The type is string. This is a key dimension.

#### **Metrics**

Physical Disk Read Bytes/sec (Superseded)

- The rate bytes are transferred from the disk during read operations.
- The type is int.
- The unit is bytes/second.

Physical Average Disk Read Queue Length

- The average number of read requests that were queued for the selected disk during the sample interval.
- The type is double.
- The unit is requests.

Physical Disk Bytes/sec (Superseded)

- The rate at which bytes are transferred to or from a disk during write or read operations.
- The type is int.
- The unit is bytes/second.

Physical % Disk Write Time

- The percentage of elapsed time that a disk drive has been busy servicing write requests. Valid values are positive integers in the range 0 to 100 (expressing a percentage) and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions.
- The type is int.

- The unit is percent.

#### Avg Disk ms/Transfer

- Time in milliseconds of the average disk transfer.
- The type is int.
- The unit is milliseconds.

#### Physical Disk Write Bytes/sec (Superseded)

- The rate bytes are transferred to the disk during write operations.
- The type is int.
- The unit is bytes/second.

#### Physical Average Disk Write Queue Length

- The average number of write requests that were queued for the selected disk during the sample interval.
- The type is double.
- The unit is requests.

#### Disk Queue Length

- The number of requests outstanding on a disk the instant the data is collected, including requests currently in service.
- The type is int.
- The unit is requests.

#### Physical Disk Read Bytes/sec

- The rate bytes are transferred from the disk during read operations. This attribute is the 64-bit version of Disk Read Bytes/sec.
- The type is double.
- The unit is bytes/second.

#### Physical % Disk Read Time

- The percentage of elapsed time a disk drive been busy servicing read requests.
- The type is int.
- The unit is percent.

#### Avg Disk Bytes/Write (Superseded)

- Average number of bytes transferred to the disk during write operations.
- The type is int.
- The unit is bytes.

#### Physical Average Disk Queue Length

- The average number of both read and write requests that were queued for the selected disk during the sample interval.
- The type is double.
- The unit is requests.

#### % Disk Idle Time

- The percentage of elapsed time that the selected disk drive is not servicing any read or write requests.
- The type is int.
- The unit is percent.

#### Avg Disk Bytes/Write

- Average number of bytes transferred to the disk during write operations. This attribute is the 64-bit version of Avg Disk Bytes/Write.
- The type is double.
- The unit is bytes.

#### Physical Avg Disk ms/Read

- Average time in milliseconds of a read of data from the disk.
- The type is int.
- The unit is milliseconds.

#### Physical Disk Transfers/sec

- The average number of read and write operations that have occurred on a disk per second.
- The type is int.
- The unit is transfers/second.

#### Avg Disk Bytes/Read

- The average number of bytes transferred from a disk during read operations. This attribute is the 64-bit version of Avg Disk Bytes/Read.
- The type is double.
- The unit is bytes.

#### Avg Disk ms/Write

- Average time in milliseconds of a write of data to the disk.
- The type is int.
- The unit is milliseconds.

#### Physical Disk Bytes/sec

- The rate at which bytes are transferred to or from a disk during write or read operations. This attribute is the 64-bit version of Disk Bytes/Sec.
- The type is double.
- The unit is bytes/second.

#### Avg Disk Bytes/Read (Superseded)

- The average number of bytes transferred from a disk per read operation.
- The type is int.
- The unit is bytes.

#### Avg Disk Bytes/Transfer

- Average number of bytes transferred to or from the disk during write or read operations. This attribute is the 64-bit version of Avg Disk Bytes/Transfer.
- The type is double.
- The unit is bytes.

#### Physical Disk Write Bytes/sec

- The rate bytes are transferred to the disk during write operations. This attribute is the 64-bit version of Disk Write Bytes/sec.
- The type is double.
- The unit is bytes/second.

#### Physical Disk Reads/sec

- The average number of read operations that have occurred on a disk per second.
- The type is int.
- The unit is reads/second.

#### Row Number

- Row Number.
- The type is int.
- The unit is row.

#### Avg Disk Bytes/Transfer (Superseded)

- Average number of bytes transferred to or from the disk during write or read operations.
- The type is int.
- The unit is bytes.

#### Physical % Disk Time

- The percentage of elapsed time a disk drive has been busy servicing read or write requests. Valid values are positive integers in the range 0 to 100 (expressing a percentage) and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions.
- The type is int.
- The unit is percent.

#### Physical Disk Writes/sec

- The average number of write operations that have occurred on a disk per second.
- The type is int.
- The unit is writes/second.

### **Component: Objects**

Information about the number of events, mutexes, processes, sections, semaphores, and threads.

#### **Dimensions**

##### Objects Parameter

- A special one for the front end to use as a column header.
- The type is string.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### Objects Value

- A special one for the front end to use as a column header.



- The type is string.

node

- The managed system name.
- The type is string.

## **Metrics**

Objects Mutexes

- The number of mutexes on a system at the time of monitoring. This is an instantaneous count, not an average. The system uses mutexes to assure that only one section of code is executing per thread.
- The type is int.
- The unit is mutexes.

Objects Sections

- The number of sections on a system at the time of monitoring. A process creates sections in memory to store data.
- The type is int.
- The unit is sections.

Objects Events

- The number of events on a system at the time of monitoring. Note that an event is any system or user action that causes notification or a log entry.
- The type is int.
- The unit is events.

Row Number

- Row Number.
- The type is int.
- The unit is row.

Objects Threads

- The number of threads on a system at the time of monitoring.
- The type is int.
- The unit is threads.

Objects Number of active processes

- The number of active processes on a system at the time of monitoring. This is an instantaneous count, not an average. A process is a running program. Note: the value 2147483647 indicates Value\_Exceeds\_Maximum and the value -2147483648 indicates Value\_Exceeds\_Minimum.
- The type is int.
- The unit is processes.

Objects Semaphores

- The number of semaphores on a system at the time of monitoring. Semaphores allow threads access to data structures that they share with other threads.
- The type is int.
- The unit is semaphores.

## Component: Log Profile Events

Information about events matching configured formats in monitored log files.

### Dimensions

Log Profile Events Timestamp

- This is the local time when the data was collected.
- The type is timestamp.

Log Profile Events Agent System Name

- This is the managed system name of the agent.
- The type is string.

Log Profile Events Class

- The Class name of the log file event, as defined in the configured format (.fmt) file.
- The type is string. This is a key dimension.

Log Profile Events CustomSlot2

- The user-defined slot from the EIF event.
- The type is string.

Log Profile Events System Name

- This is the managed system name of the agent.
- The type is string.

Log Profile Events CustomSlot3

- The user-defined slot from the EIF event.
- The type is string.

Log Profile Events CustomSlot1

- The user-defined slot from the EIF event.
- The type is string.

Log Profile Events Logkey

- The name of the log in which the matching record was found.
- The type is string. This is a key dimension.

Log Profile Events CustomSlot6

- The user-defined slot from the EIF event.
- The type is string.

Log Profile Events CustomSlot7

- The user-defined slot from the EIF event.
- The type is string.

Log Profile Events CustomSlot4

- The user-defined slot from the EIF event.
- The type is string.

#### Log Profile Events CustomSlot5

- The user-defined slot from the EIF event.
- The type is string.

#### Log Profile Events Event Type

- A flag indicating whether the current event is a flood control summary event.
- The type is int.

#### Log Profile Events EventId

- The event ID per subnode and log file name.
- The type is double.

#### Log Profile Events CustomSlot8

- The user-defined slot from the EIF event.
- The type is string.

#### Log Profile Events CustomSlot9

- The user-defined slot from the EIF event.
- The type is string.

#### Log Profile Events CustomInteger1

- The user-defined slot with integral type from the EIF event.
- The type is double.

#### Log Profile Events CustomInteger3

- The user-defined slot with integral type from the EIF event.
- The type is double.

#### Log Profile Events CustomInteger2

- The user-defined slot with integral type from the EIF event.
- The type is double.

#### Log Profile Events Logname

- The name of the log in which the matching record was found.
- The type is string.

#### Log Profile Events CustomSlot10

- The user-defined slot from the EIF event.
- The type is string.

#### Log Profile Events RemoteHost

- The remote host on which the event originated if using ssh/remote, blank for local system.
- The type is string.

#### Log Profile Events EIFEvent

- The contents of the log file event in EIF format.
- The type is string.

## Metrics

Log Profile Events Occurrence Count

- The number of times this event occurred over the current flood control summary interval.
- The type is int.
- The unit is count.

Log Profile Events msg

- The contents of the msg slot from the EIF event.
- The type is string.
- The unit is string.

## Component: Log File Status

Information that reflects the status of log files this agent is monitoring.

### Dimensions

Log File Status System Name

- This is the managed system name of the agent.
- The type is string.

Log File Status Subnode Resource Name

- The Resource Name of the subnode agent.
- The type is string.

Log File Status Table Name

- The name of the table in which this log is monitored.
- The type is string. This is a key dimension.

Log File Status File Name

- The full path name of the monitored file. The file name is null if the row represents a file pattern to be matched during the next scan.
- The type is string. This is a key dimension.

Log File Status Logkey

- The name of the log in which the matching record was found. Used to map the row with the LogfileProfileEvents group.
- The type is string. This is a key dimension.

Log File Status Current File Position

- The current position in bytes into the monitored file. Data up to this has been processed, after that has not. Not applicable to pipes.
- The type is double.

Log File Status File Status

- The current status of this file.
- The type is int.

Log File Status Last Modification Time

- The time when the monitored file was last written to. Not applicable to pipes.

- The type is timestamp.

#### Log File Status RegEx Pattern

- The regular expression pattern (if any) that caused this file to be monitored. The regular expression may be the same as the File Name attribute if the regular expression contains no meta characters or if the dotmeta character actually matches a dot. If the regular expression is different than the non-null File Name attribute, then the regular expression pattern was used to match the filename.
- The type is string. This is a key dimension.

#### Log File Status Codepage

- The language codepage of the monitored file.
- The type is string.

#### Log File Status Timestamp

- This is the local time when the data was collected.
- The type is timestamp.

#### Log File Status File Type

- The type of this file (regular file or pipe).
- The type is int.

#### Log File Status Remote Host

- The remote host name where the monitored file resides. The host name will be N/A if the row represents a local file.
- The type is string. This is a key dimension.

### Metrics

#### Log File Status Num Records Not Matched

- The number of processed records from this log which did not match any of the specified patterns.
- The type is int.
- The unit is records.

#### Log File Status Current File Size

- The current size of the monitored file. Not applicable to pipes.
- The type is double.
- The unit is bytes.

#### Log File Status Num Records Matched

- The number of processed records from this log which matched one of the specified patterns.
- The type is int.
- The unit is records.

#### Log File Status Num Records Processed

- The number of records processed from this log since agent start (including ones that are not matches/events).
- The type is int.

- The unit is records.

### **Component: Gopher Service**

Information about traffic and connection activity for a Gopher server, such as the current connections, the bytes received per second, and the total non-anonymous users connected.

#### **Dimensions**

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

Gopher Service Parameter

- A special one for the front end to use as a column header.
- The type is string.

Gopher Service Value

- A special one for the front end to use as a column header.
- The type is string.

node

- The managed system name.
- The type is string.

#### **Metrics**

Gopher Service Aborted Connections

- The number of aborted connections.
- The type is int.
- The unit is connections.

Gopher Service Total Anonymous Users

- The total number of anonymous users that have ever connected to the Gopher server.
- The type is int.
- The unit is users.

Gopher Service Files sent by Gopher Server

- The total number of files sent by the Gopher server.
- The type is int.
- The unit is files.

Gopher Service HTTP Service Bytes Total/sec

- The total number of KBs flowing through the Gopher server per second. This includes both incoming and outgoing bytes. This number is a good indicator of how heavily your Gopher server is loaded.
- The type is int.
- The unit is kilobytes/second.

Gopher Service Bytes Sent/sec

- The rate that bytes are sent by the Gopher server.

- The type is int.
- The unit is bytes/second.

#### Gopher Service Maximum Anonymous Users

- The maximum number of anonymous users simultaneously connected to the Gopher server.
- The type is int.
- The unit is users.

#### Gopher Svc Bytes Received/sec

- The rate that bytes are received by the Gopher server.
- The type is int.
- The unit is bytes/second.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### Gopher Service Gopher Plus Requests

- The number of Gopher Plus requests received by the Gopher server.
- The type is int.
- The unit is requests.

#### Gopher Service Directory Listings Sent

- The total number of directory listings sent by the Gopher server.
- The type is int.
- The unit is listings.

#### Gopher Service Logon Attempts

- The number of logon attempts that have been made by the Gopher server.
- The type is int.
- The unit is attempts.

#### Gopher Service Connections in Error

- The number of connections that had errors when processed by the Gopher Server.
- The type is int.
- The unit is connections.

#### Gopher Service Maximum NonAnonymous Users

- The maximum number of non-anonymous users simultaneously connected to the Gopher server.
- The type is int.
- The unit is users.

#### Gopher Service Total NonAnonymous Users

- The total number of non-anonymous users that have ever connected to the Gopher server.
- The type is int.
- The unit is users.

#### Gopher Service Connection Attempts

- The number of connections that had errors when processed by the Gopher server.
- The type is int.
- The unit is connections.

#### Gopher Service Current Anonymous Users

- The number of anonymous users currently connected to the Gopher server.
- The type is int.
- The unit is users.

#### Gopher Service Current NonAnonymous Users

- The number of non-anonymous users currently connected to the Gopher server.
- The type is int.
- The unit is users.

#### Gopher Service Searches Sent

- The total number of searches performed by the Gopher server.
- The type is int.
- The unit is searches.

#### Gopher Service Current Connections

- The current number of connections to the Gopher server.
- The type is int.
- The unit is connections.

#### Gopher Service Maximum Connections

- The number of connection attempts that have been made to the Gopher server.
- The type is int.
- The unit is connections.

### **Component: DNS Dynamic Update**

Information about DNS (Domain Name Server) server activity and performance.

#### **Dimensions**

##### DNS Dynamic DNS Value

- A special one for the front end to use as a column header.
- The type is string.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### DNS Dynamic Node



- The managed system name.
- The type is string.

#### DNS Dynamic DNS Parameter

- A special one for the front end to use as a column header.
- The type is string.

### Metrics

#### DNS Dynamic Update Rejected

- The total number of dynamic updates rejected by the DNS server.
- The type is int.
- The unit is requests.

#### DNS Dynamic DNS Secure Update Failure

- The total number of secure updates failed of the DNS server.
- The type is int.
- The unit is errors.

#### DNS Dynamic Update NoOperation

- The total number of No-operation/Empty dynamic update requests received by the DNS server.
- The type is int.
- The unit is requests.

#### DNS Dynamic DNS Secure Update Received/sec

- The average number of secure update requests received by the DNS server in each second.
- The type is int.
- The unit is requests/seconds.

#### DNS Dynamic Update Received/sec

- The average number of dynamic update requests received by the DNS server in each second.
- The type is int.
- The unit is requests/second.

#### DNS Dynamic Update Written to Database/sec

- The average number of dynamic updates written to the database by the DNS server in each second.
- The type is int.
- The unit is requests/second.

#### DNS Dynamic DNS Secure Update Received

- The total number of secure update requests received by the DNS server.
- The type is int.
- The unit is requests.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### DNS Dynamic Update Written to Database

- The total number of dynamic updates written to the database by the DNS server.
- The type is int.
- The unit is requests.

#### DNS Dynamic Update Received

- The total number of dynamic update requests received by the DNS server.
- The type is int.
- The unit is requests.

#### DNS Dynamic Update TimeOuts

- The total number of dynamic update timeouts of the DNS server.
- The type is int.
- The unit is timeouts.

#### DNS Dynamic Update Queued

- The total number of dynamic updates queued by the DNS server.
- The type is int.
- The unit is requests.

#### DNS Dynamic Update NoOperation/sec

- The average number of No-operation/Empty dynamic update requests received by the DNS server in each second.
- The type is int.
- The unit is requests/second.

### **Component: RAS Total**

Information about Total Remote Access Service activity.

#### **Dimensions**

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### node

- The managed system name.
- The type is string.

##### RAS Total Value

- A special one for the front end to use as a column header.
- The type is string.

#### RAS Total Parameter

- A special one for the front end to use as a column header.
- The type is string.

### **Metrics**

#### RAS Total Frames Transmitted

- The total number of data frames transmitted for this connection.
- The type is int.
- The unit is frames.

#### RAS Total Buffer Overrun Errors

- The total number of buffer overrun errors for this connection. Buffer overrun errors occur when the software cannot handle the rate at which data is received.
- The type is int.
- The unit is errors.

#### RAS Total Frames Received

- The total number of data frames received for this connection.
- The type is int.
- The unit is frames.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### RAS Total Percent Compression Out

- The compression ratio for bytes being transmitted.
- The type is int.
- The unit is percent.

#### RAS Total Total Connections

- The total number of Remote Access connections.
- The type is int.
- The unit is connections.

#### Ras Total Bytes Transmitted

- The total number of bytes transmitted for this connection.
- The type is int.
- The unit is bytes.

#### RAS Total Total Errors

- The total number of CRC, Timeout, Serial Overrun, Alignment, and Buffer Overrun errors for this connection.
- The type is int.

- The unit is errors.

#### RAS Total Bytes Received/sec

- The number of bytes received per second for this connection.
- The type is int.
- The unit is bytes/second.

#### RAS Total CRC Errors

- The total number of CRC errors for this connection. CRC errors occur when the frame received contains erroneous data.
- The type is int.
- The unit is errors.

#### RAS Total Percent Compression In

- The compression ratio for bytes being received.
- The type is int.
- The unit is percent.

#### RAS Total Serial Overrun Errors

- The total number of serial overrun errors for this connection. Serial overrun errors occur when the hardware cannot handle the rate at which data is received.
- The type is int.
- The unit is errors.

#### RAS Total Bytes Received

- The total number of bytes received for this connection.
- The type is int.
- The unit is bytes.

#### RAS Total Frames Transmitted/sec

- The number of data frames transmitted per second for this connection.
- The type is int.
- The unit is frames/second.

#### RAS Total Total Errors/sec

- The total number of CRC, Timeout, Serial Overrun, Alignment, and Buffer Overrun errors per second.
- The type is int.
- The unit is errors/second.

#### RAS Total Alignment Errors

- The total number of alignment errors for this connection. Alignment errors occur when a byte received is different from the byte expected.
- The type is int.
- The unit is errors.

#### Ras Total Bytes Transmitted/sec

- The number of bytes transmitted per second for this connection.

- The type is int.
- The unit is bytes/second.

#### RAS Total Timeout Errors

- The total number of timeout errors for this connection. Timeout errors occur when an expected response is not received in time.
- The type is int.
- The unit is errors.

#### RAS Total Frames Received/sec

- The number of data frames received per second for this connection.
- The type is int.
- The unit is frames/second.

### **Component: Active Tasks running on VCenter Server**

Use the Web Service data set to create situations to monitor traffic and connection activity for a web server. Web Service is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

node

- The managed system name.
- The type is string.

Web Service Web Site Instance Name

- Name of web site. Valid format is a text string of up to 64 characters.
- The type is string. This is a key dimension.

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### **Metrics**

Web Service Bytes Received/sec (Superseded)

- The rate that data bytes are received by the Web service.
- The type is int.
- The unit is bytes/second.

Web Service Total Anonymous Users

- The total number of users who established an anonymous connection with the Web service (counted since service startup).
- The type is int.
- The unit is users.

Web Service Bytes Sent/sec (Superseded)

- The rate that data bytes are sent by the Web service.
- The type is int.
- The unit is bytes/second.

#### Web Service Total Head Requests

- The number of HTTP requests using the HEAD method (counted since service startup). Head requests generally indicate a client is querying the state of a document they already have to see if it needs to be refreshed.
- The type is int.
- The unit is requests.

#### Web Service ISAPI Extension Requests/sec

- The rate of ISAPI Extension requests that are simultaneously being processed by the Web service.
- The type is int.
- The unit is requests/second.

#### Web Service Total Get Requests

- The number of HTTP requests using the GET method (counted since service startup). Get requests are generally used for basic file retrievals or image maps, though they can be used with forms.
- The type is int.
- The unit is requests.

#### Web Service Current Anonymous Users

- The number of users who currently have an anonymous connection using the Web service.
- The type is int.
- The unit is users.

#### Web Service Post Requests/sec

- The rate HTTP requests using the POST method are made. Post requests are generally used for forms or gateway requests.
- The type is int.
- The unit is requests/second.

#### Web Service Total Not Found Errors

- The number of requests that could not be satisfied by the server because the requested document could not be found. These are generally reported as an HTTP 404 error code to the client. The count is the total since service startup.
- The type is int.
- The unit is errors.

#### Web Service Maximum CGI Requests

- Maximum number of CGI requests simultaneously processed by the Web service.
- The type is int.
- The unit is requests.

#### Web Service Current ISAPI Extension Requests

- The current number of Extension requests that are simultaneously being processed by the Web service.
- The type is int.
- The unit is requests.

#### Web Service Total Rejected Async I/O Requests

- Total requests rejected due to bandwidth throttling settings (counted since service startup).
- The type is int.
- The unit is requests.

#### Web Service Maximum Connections

- The maximum number of simultaneous connections established with the Web service.
- The type is int.
- The unit is connections.

#### Web Service Logon Attempts/sec

- The rate that logons using the Web service are being attempted.
- The type is int.
- The unit is attempts/second.

#### Web Service Anonymous Users/sec

- The rate users are making anonymous connections using the Web service.
- The type is int.
- The unit is users/second.

#### Web Service Files/sec

- The rate files are transferred, that is, sent and received by the Web service.
- The type is int.
- The unit is files/second.

#### Web Service Total Files Received

- The total number of files received by the Web service (counted since service startup).
- The type is int.
- The unit is files.

#### Web Service Maximum Anonymous Users

- The maximum number of users who established concurrent anonymous connections using the Web service (counted since service startup).
- The type is int.
- The unit is users.

#### Web Service Head Requests/sec

- The rate HTTP requests using the HEAD method are made. Head requests generally indicate a client is querying the state of a document they already have to see if it needs to be refreshed.
- The type is int.
- The unit is requests/second.

#### Web Service Total Delete Requests

- Total Delete Requests is the number of HTTP requests using the DELETE method (counted since service startup). Delete requests are generally used for file removals.
- The type is int.
- The unit is requests.

#### Web Service Get Requests/sec

- The rate HTTP requests using the GET method are made. Get requests are generally used for basic file retrievals or image maps, though they can be used with forms.
- The type is int.
- The unit is requests/second.

#### Web Service Other Request Methods/sec

- The rate HTTP requests are made that do not use the GET, POST, PUT, DELETE, TRACE or HEAD methods. These might include LINK or other methods supported by gateway applications.
- The type is int.
- The unit is methods/second.

#### Web Service Bytes Sent/sec

- The rate that data bytes are sent by the Web service. This attribute is the 64-bit version of Bytes\_Sent/sec.
- The type is double.
- The unit is bytes/second.

#### Web Service Delete Requests/sec

- The rate HTTP requests using the DELETE method are made. Delete requests are generally used for file removals.
- The type is int.
- The unit is requests/second.

#### Web Service Total Allowed Async I/O Requests

- Total requests allowed by bandwidth throttling settings (counted since service startup).
- The type is int.
- The unit is requests.

#### Web Service Total Method Requests

- The number of HTTP GET, POST, PUT, DELETE, TRACE, HEAD and other method requests (counted since service startup).
- The type is int.
- The unit is requests.

#### Web Service Current NonAnonymous Users

- The number of users who currently have a non-anonymous connection using the Web service. Note: the value 2147483647 indicates Value\_Exceeds\_Maximum and the value -2147483648 indicates Value\_Exceeds\_Minimum.
- The type is int.
- The unit is users.

#### Web Service Current Blocked Async I/O Requests

- Current requests temporarily blocked due to bandwidth throttling settings.
- The type is int.
- The unit is requests.

#### Web Service Total Other Request Methods



- The number of HTTP requests that are not GET, POST, PUT, DELETE, TRACE or HEAD methods (counted since service startup). These might include LINK or other methods supported by gateway applications.
- The type is int.
- The unit is requests.

#### Web Service Bytes Received/sec

- The rate that data bytes are received by the Web service. This attribute is the 64-bit version of Bytes\_Received/sec.
- The type is double.
- The unit is bytes/second.

#### Web Service Total Connection Attempts

- The number of connections that have been attempted using the Web service (counted since service startup).
- The type is int.
- The unit is attempts.

#### Web Service Total NonAnonymous Users

- The total number of users who established a non-anonymous connection with the Web service (counted since service startup).
- The type is int.
- The unit is users.

#### Web Service Bytes Total/sec (Superseded)

- The sum of Bytes Sent/sec and Bytes Received/sec. This is the total rate of bytes transferred by the Web service. Note: the value 2147483647 indicates Value\_Exceeds\_Maximum and the value -2147483648 indicates Value\_Exceeds\_Minimum.
- The type is int.
- The unit is bytes/second.

#### Web Service Total Blocked Async I/O Requests

- Total requests temporarily blocked due to bandwidth throttling settings (counted since service startup).
- The type is int.
- The unit is requests.

#### Web Service Total Method Requests/sec

- The rate HTTP requests using GET, POST, PUT, DELETE, TRACE or HEAD methods are made.
- The type is int.
- The unit is requests/second.

#### Web Service Connection Attempts/sec

- The rate that connections using the Web service are being attempted.
- The type is int.
- The unit is attempts.

#### Web Service Files Received/sec

- The rate files are received by the Web service.
- The type is int.
- The unit is files/second.

#### Web Service Put Requests/sec

- The rate HTTP requests using the PUT method are made.
- The type is int.
- The unit is requests/second.

#### Web Service Not Found Errors/sec

- The rate of errors due to requests that could not be satisfied by the server because the requested document could not be found. These are generally reported as an HTTP 404 error code to the client.
- The type is int.
- The unit is errors/second.

#### Web Service Current Connections

- The current number of connections established with the Web service.
- The type is int.
- The unit is connections.

#### Web Service Total Files Sent

- The total number of files sent by the Web service (counted since service startup).
- The type is int.
- The unit is files.

#### Web Service Current CGI Requests

- Current number of CGI requests that are simultaneously being processed by the Web service.
- The type is int.
- The unit is requests.

#### Web Service Total Logon Attempts

- The number of logons that have been attempted using the Web service (counted since service startup).
- The type is int.
- The unit is attempts.

#### Web Service Total CGI Requests

- Custom gateway executables (exe) the administrator can install to add forms processing or other dynamic data sources. CGI requests spawn a process on the server which can be a large drain on server resources. The count is the total since service startup.
- The type is int.
- The unit is requests.

#### Web Service CGI Requests/sec

- The rate of CGI requests that are simultaneously being processed by the Web service.
- The type is int.
- The unit is requests/second.

#### Web Service Total Put Requests

- The number of HTTP requests using the PUT method (counted since service startup).
- The type is int.
- The unit is requests.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### Web Service Maximum ISAPI Extension Requests

- The maximum number of Extension requests simultaneously processed by the Web service.
- The type is int.
- The unit is requests.

#### Web Service Total Files Transferred

- The sum of Files Sent and Files Received. This is the total number of files transferred by the Web service (counted since service startup).
- The type is int.
- The unit is files.

#### Web Service Total Trace Requests

- The number of HTTP requests using the TRACE method (counted since service startup). Trace requests allow the client to see what is being received at the end of the request chain and use the information for diagnostic purposes.
- The type is int.
- The unit is requests.

#### Web Service Total ISAPI Extension Requests

- Custom gateway Dynamic Link Libraries (dll) the administrator can install to add forms processing or other dynamic data sources. Unlike CGI requests, ISAPI requests are simple calls to a DLL library routine, thus they are better suited to high performance gateway applications. The count is the total since service startup.
- The type is int.
- The unit is requests.

#### Web Service System Code Resident Bytes

- System Code Resident Bytes. Note that this attribute is not available on systems running Windows 2000 with IIS 5.0).
- The type is int.
- The unit is bytes.

#### Web Service Measured Async I/O Bandwidth Usage

- Measured bandwidth of asynchronous I/O averaged over a minute.
- The type is int.
- The unit is ?.

#### Web Service NonAnonymous Users/sec

- The rate users are making non-anonymous connections using the Web service.
- The type is int.
- The unit is users/second.

#### Web Service Total Post Requests

- The number of HTTP requests using the POST method (counted since service startup). Post requests are generally used for forms or gateway requests.
- The type is int.
- The unit is requests.

#### Web Service Files Sent/sec

- The rate files are sent by the Web service.
- The type is int.
- The unit is files/second.

#### Web Service Maximum NonAnonymous Users

- The maximum number of users who established concurrent non-anonymous connections using the Web service (counted since service startup).
- The type is int.
- The unit is users.

#### Web Service Bytes Total/sec

- The sum of Bytes Sent/sec and Bytes Received/sec. This is the total rate of bytes transferred by the Web service. This attribute is the 64-bit version of Bytes Total/sec.
- The type is double.
- The unit is bytes/second.

### **Component: Active Tasks running on VCenter Server**

Use the Server Work Queue data set to monitor information about server work queue throughput, work items in the queue, and threads servicing the queue. This data set is superseded. There is a new data set with the same name that replaces it.

#### **Dimensions**

node

- The managed system name.
- The type is string.

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

Server Work Queue Name

- Instance Name. Valid format is a text string of up to 64 characters.
- The type is string. This is a key dimension.

#### **Metrics**

Server Work Queue Available Threads

- The number of server threads on this CPU not currently working on requests from a client. The server dynamically adjusts the number of threads to maximize server performance.
- The type is int.
- The unit is threads.

#### Server Work Queue Current Clients

- The instantaneous count of the clients being serviced by this CPU. The server actively balances the client load across all of the CPU's in the system.
- The type is int.
- The unit is clients.

#### Work Item Shortages

- The number of times a request waited for an available workitem from the pool. Every request from a client is represented in the server as a 'work item,' and the server maintains a pool of available work items per CPU to speed processing. A sustained value greater than zero indicates the need to increase the 'MaxWorkItems' registry value for the Server service. This value is always 0 in the Blocking Queue instance.
- The type is int.
- The unit is waits.

#### Write Operations/sec

- The rate the server is performing file write operations for the clients on this CPU. This value is a measure of how busy the Server is. This value is always 0 in the Blocking Queue instance.
- The type is double.
- The unit is operations/second.

#### Server Work Queue Available Work Items

- The instantaneous number of available work items for this CPU. A sustained near-zero value indicates the need to increase the MinFreeWorkItems registry value for the Server service. This value is always 0 in the Blocking Queue instance.
- The type is int.
- The unit is workitems.

#### Server Work Queue Total Bytes/sec

- The rate the Server is reading and writing data to and from the files for the clients on this CPU. This value is a measure of how busy the Server is.
- The type is double.
- The unit is bytes/second.

#### Server Work Queue Bytes Sent/sec

- The rate at which the Server is sending bytes to the network clients on this CPU. This value is a measure of how busy the Server is.
- The type is double.
- The unit is bytes/second.

#### Total Operations/sec

- The rate the Server is performing file read and file write operations for the clients on this CPU. This value is a measure of how busy the Server is. This value is always 0 in the Blocking Queue instance.

- The type is double.
- The unit is operations.

#### Write Bytes/sec

- The rate the server is writing data to files for the clients on this CPU. This value is a measure of how busy the Server is.
- The type is double.
- The unit is bytes/second.

#### Server Work Queue Svc Work Queue Bytes Received/sec

- The rate at which the Server is receiving bytes from the network clients on this CPU. This value is a measure of how busy the Server is.
- The type is double.
- The unit is bytes/second.

#### Server Work Queue Bytes Transferred/sec

- The rate at which the Server is sending and receiving bytes with the network clients on this CPU. This value is a measure of how busy the Server is.
- The type is double.
- The unit is bytes/second.

#### Server Work Queue Active Threads

- The number of threads currently working on a request from the server client for this CPU. The system keeps this number as low as possible to minimize unnecessary context switching. This is an instantaneous count for the CPU, not an average over time.
- The type is int.
- The unit is threads.

#### Server Work Queue Read Operations/sec

- The rate the server is performing file read operations for the clients on this CPU. This value is a measure of how busy the Server is. This value is always 0 in the Blocking Queue instance.
- The type is double.
- The unit is operations/second.

#### Server Work Queue Read Bytes/sec

- The rate the server is reading data from files for the clients on this CPU. This value is a measure of how busy the Server is.
- The type is double.
- The unit is bytes/second.

#### Server Work Queue Context Blocks Queued/sec

- The rate at which work context blocks had to be placed on the FSP queue of the server to await server action.
- The type is int.
- The unit is blocks/second.

#### Server Work Queue Length

- The current length of the server work queue for this CPU. A sustained queue length greater than four might indicate processor congestion. This is an instantaneous count, not an average over time.
- The type is int.
- The unit is queues.

#### Server Work Queue Borrowed Work Items

- The number of borrowed work items. Every request from a client is represented in the server as a 'work item,' and the server maintains a pool of available work items per CPU to speed processing. When a CPU runs out of work items, it borrows a free work item from another CPU. An increasing value of this running counter might indicate the need to increase the 'MaxWorkItems' or 'MinFreeWorkItems' registry values for the Server service. This value is always 0 in the Blocking Queue instance.
- The type is int.
- The unit is workitems.

### **Component: DNS Memory**

Information about DNS (Domain Name Server) server activity and performance.

#### **Dimensions**

##### DNS Memory node

- The managed system name.
- The type is string.

##### DNS Memory DNS Value

- A special one for the front end to use as a column header.
- The type is string.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### DNS Memory DNS Parameter

- A special one for the front end to use as a column header.
- The type is string.

#### **Metrics**

##### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

##### DNS Memory DNS TCP Message Memory

- The total TCP message memory used by DNS server.
- The type is int.
- The unit is bytes.

#### DNS Memory DNS Caching Memory

- The total caching memory used by DNS server.
- The type is int.
- The unit is bytes.

#### DNS Memory DNS UDP Message Memory

- The total UDP message memory used by DNS server.
- The type is int.
- The unit is bytes.

#### DNS Memory DNS Database Node Memory

- The total database node memory used by DNS server.
- The type is int.
- The unit is ?.

#### DNS Memory DNS Record Flow Memory

- The total record flow memory used by DNS server.
- The type is int.
- The unit is ?.

#### DNS Memory DNS Nostat Memory

- The total Nostat memory used by DNS server.
- The type is int.
- The unit is ?.

### **Component: Active Tasks running on VCenter Server**

Use the Network Port data set to monitor connection information about network ports.

#### **Dimensions**

##### Local Port

- The local port number.
- The type is int. This is a key dimension.

##### Remote Port

- The remote port number. Note: 0 indicates Unavailable.
- The type is int.

##### State

- The port state. For sessions that are established with a protocol of UDP, the remote port and state attributes will have a value of Unavailable.
- The type is int.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### Local Host Name



- The host name of the local host. This name can either be the simple host name or the fully qualified host name.
- The type is string.

#### Remote Host Name

- The host name of the remote host when the protocol is TCP. This value is blank if the protocol is UDP. This name can either be the simple host name or the fully qualified host name.
- The type is string.

#### Protocol

- The port protocol.
- The type is string. This is a key dimension.

#### Local Port Name

- The local port name.
- The type is string.

#### Remote Host IP Address

- The IP address of the remote host when the protocol is TCP. This value is blank if the protocol is UDP.
- The type is string.

#### Local Host IP Address

- The IP address of the local host.
- The type is string.

#### node

- The managed system name.
- The type is string.

### **Component: Active Tasks running on VCenter Server**

Use the Logical Disk data set to create situations that monitor information about disk drive partitions that have been assigned a drive letter. Logical disk is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### Logical Disk Value

- A special one for the front end to use as a column header.
- The type is string.

##### node

- The managed system name.
- The type is string.

##### Logical Disk Parameter

- A special one for the front end to use as a column header.
- The type is string.

##### Logical Disk Name

- The name of a logical disk.
- The type is string. This is a key dimension.

#### Logical Disk Size

- The size of the logical disk, in MBs. 1 MB = 1,048,576 bytes. Note: -1 indicates Unavailable.
- The type is int.

#### Logical Disk Name (Long)

- The long name of the logical disk.
- The type is string.

#### Logical Disk Physical Disk Number

- The number of the physical disk that contains this logical disk. This value is only provided for drives assigned as drive letters. If the logical disk is a mounted volume, this value will be "Mounted". Note: Mnt = Mounted.
- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

### Metrics

#### Logical Disk Percent Volume Free Space

- The percentage of the volume that is free space.
- The type is int.
- The unit is percent.

#### Logical Disk Bytes/sec

- The rate at which the system has transferred bytes to and from a logical disk during write or read operations. This attribute is the 64-bit version of Disk Bytes/ sec.
- The type is double.
- The unit is bytes/second.

#### Logical Disk % Disk Write Time

- The percentage of elapsed time that a logical disk drive has been busy servicing write requests.
- The type is int.
- The unit is percent.

#### Logical Disk Bytes/sec (Superseded)

- The rate at which the system has transferred bytes to and from a logical disk during write or read operations.
- The type is int.
- The unit is bytes/second.

#### Logical Disk Read Bytes/sec (Superseded)

- The rate at which the system has transferred bytes from a logical disk during read operations.
- The type is int.
- The unit is bytes/second.

#### Logical Disk Percent Used

- The percentage of the volume that is used.
- The type is int.
- The unit is percent.

#### Logical Disk Read Bytes/sec

- The rate at which the system has transferred bytes from a logical disk during read operations. This attribute is the 64-bit version of Disk Read Bytes/sec.
- The type is double.
- The unit is bytes/second.

#### Row Number

- Row Number.
- The type is int.
- The unit is row.

#### Logical Disk Average Disk Queue Length

- The average number of both read and write requests that were queued for the selected disk during the sample interval.
- The type is double.
- The unit is requests.

#### Logical Disk Disk Queue Length (Requests)

- The number of requests outstanding on a logical disk. This number includes requests in service when the data is collected.
- The type is int.
- The unit is requests.

#### Logical Disk Writes/sec

- The rate of write operations to a logical disk.
- The type is int.
- The unit is writes/second.

#### Logical Disk Unallocated Space

- The number of MBs of unallocated space on a logical drive. Note: 1 MB = 1,048,576 bytes.
- The type is int.
- The unit is megabytes.

#### Logical Disk Transfers/sec

- The rate of read and write operations on a logical disk.
- The type is int.
- The unit is transfers/second.

#### Logical Disk Write Bytes/sec (Superseded)

- The rate at which the system has transferred bytes to a logical disk during write operations.
- The type is int.
- The unit is bytes/second.

#### Logical Disk % Disk Read Time

- The percentage of elapsed time a logical disk has been busy servicing read requests.
- The type is int.
- The unit is percent.

#### Logical Disk Reads/sec

- The rate of read operations from a logical disk.
- The type is int.
- The unit is reads/second.

#### Logical Disk Average Disk Read Queue Length

- The average number of read requests that were queued for the selected disk during the sample interval.
- The type is double.
- The unit is requests.

#### Logical Disk Average Disk Write Queue Length

- The average number of write requests that were queued for the selected disk during the sample interval.
- The type is double.
- The unit is requests.

#### Logical Disk % Disk Time

- The percentage of elapsed time that a logical disk has been busy servicing read and write requests. Valid values are positive integers in the range 0 to 100 (expressing a percentage) and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions.
- The type is int.
- The unit is percent.

#### Logical Disk Avg Disk ms/Read

- The average amount of time for a read of data from a logical disk.
- The type is int.
- The unit is milliseconds.

#### Logical Disk Write Bytes/sec

- The rate at which the system has transferred bytes to a logical disk during write operations. This attribute is the 64-bit version of Disk Write Bytes/sec.
- The type is double.
- The unit is bytes/second.

### **Component: Cache Activity**

Information about cache activity, such as the frequency of reads from cache pages, the percentage of cache copy requests that were successful, and the number of pages the cache has flushed to disk. Cache is a single-instance data set.

#### **Dimensions**

##### Cache Value

- A special one for the front end to use as a column header.

- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### Cache Column Header

- A special one for the front end to use as a column header.
- The type is string.

#### node

- The managed system name.
- The type is string.

### **Metrics**

#### Cache Copy Read Hits Dynamic Average

- A running average of the Copy Read Hits % attribute.
- The type is int.
- The unit is percent.

#### Cache Read Aheads/sec

- The frequency of cache reads where the cache detects sequential access to a file. The read aheads permit the data to be transferred in larger blocks than those being requested by the application, reducing the overhead per access.
- The type is int.
- The unit is reads/second.

#### Cache Async Pin Reads/sec

- The frequency of reading data in to the cache prior to writing the data back to disk. Pages read in this fashion are pinned in memory at the completion of the read. The file system regains control immediately even if the disk must be accessed to retrieve the page. While pinned, the physical address for the page is not altered.
- The type is int.
- The unit is reads/second.

#### Cache Sync MDL Reads/sec

- The frequency of reads from cache pages that use a Memory Descriptor List (MDL) to access the pages. The MDL contains the physical address of each page in the transfer, thus permitting Direct Memory Access (DMA) of the pages. If the accessed page(s) are not in main memory, the caller waits for the pages to fault in from the disk.
- The type is int.
- The unit is reads/second.

#### Cache Data Map Hits %

- The percentage of data maps in the cache that could be resolved without having to retrieve a page from the disk, that is, the page was already in physical memory.
- The type is int.
- The unit is percent.

#### Cache Async MDL Reads/sec

- The frequency of reads from cache pages using a Memory Descriptor List (MDL) to access the pages. The MDL contains the physical address of each page in the transfer, thus permitting Direct Memory Access (DMA) of the pages. If the accessed page(s) are not in main memory, the calling application program does not wait for the pages to fault in from disk.
- The type is int.
- The unit is pages/second.

#### Cache Sync Pin Reads/sec

- The frequency of reading data in to the cache preparatory to writing the data back to disk. Pages read in this fashion are pinned in memory at the completion of the read. The file system does not regain control until the page is pinned in the cache, in particular if the disk must be accessed to retrieve the page. While pinned, the physical address for a page in the cache is not altered.
- The type is int.
- The unit is pins/second.

#### Cache Lazy Write Pages/sec

- The frequency with which the Lazy Write thread for the cache writes to disk. Lazy Writing is the process of updating the disk after the page has been changed in memory, so the application making the change to the file does not have to wait for the disk write to complete before proceeding. More than one page can be transferred on a single disk write operation.
- The type is int.
- The unit is writes/sec.

#### Cache MDL Read Hits %

- The percentage of cache Memory Descriptor List (MDL) read requests that hit the cache, that is, that did not require disk accesses to provide memory access to the page(s) in the cache.
- The type is int.
- The unit is percent.

#### Cache Async Fast Reads/sec

- The frequency of reads from cache pages that bypass the installed file system and retrieve the data directly from the cache. Normally, file I/O requests invoke the appropriate file system to retrieve data from a file, but this path permits direct retrieval of cache data without file system involvement if the data is in the cache. Even if the data is not in the cache, one invocation of the file system is avoided. If the data is not in the cache, the request (application program call) gets control immediately.
- The type is int.
- The unit is pages/second.

#### Cache Data Maps/sec

- The number of times per second that a file system, such as NTFS or HPFS, maps a page of a file in to the cache to read the page.
- The type is int.
- The unit is maps/second.

#### Cache Fast Reads/sec

- The frequency of reads from cache pages that bypass the installed file system and retrieve the data directly from the cache. Normally, file I/O requests invoke the appropriate file system to retrieve data from a file, but this path permits direct retrieval of cache data without file system

involvement if the data is in the cache. Even if the data is not in the cache, one invocation of the file system is avoided.

- The type is int.
- The unit is reads/second.

#### Cache Pin Read Hits %

- The percentage of cache pin read requests that hit the cache, that is, that did not require a disk read to provide access to the page in the cache. While pinned, the physical address for the page in the cache is not altered. The LAN Redirector uses this method for retrieving cache information, as does the LAN Server for small transfers. This is usually the method used by the disk file systems as well.
- The type is int.
- The unit is percent.

#### Cache Copy Reads/sec

- The frequency of reads from cache pages that involve a memory copy of the data from the cache to the buffer for the application. The LAN Redirector uses this method to retrieve cache information, as does the LAN Server for small transfers. This is a method used by the disk file systems as well.
- The type is int.
- The unit is reads/second.

#### Cache Pin Read Hits Dynamic Average

- A running average of the Pin Read Hits % attribute.
- The type is int.
- The unit is percent.

#### Cache Fast Read Resource Misses/sec

- The frequency of cache misses necessitated by the lack of available resources to satisfy the request.
- The type is int.
- The unit is misses/second.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### Cache Sync Data Maps/sec

- The frequency with which a file system, such as NTFS or HPFS, maps a page of a file in to the cache to read the page, and wishes to wait for the cache to retrieve the page if it is not in main memory.
- The type is int.
- The unit is maps/second.

#### Cache Data Flush Pages/sec

- The number of pages the cache has flushed to disk as a result of a request to flush or to satisfy a write-through file write request. More than one page can be transferred on each flush operation.
- The type is int.
- The unit is flushes/second.

#### Cache Lazy Write Flushes/sec

- The frequency with which the Lazy Write thread flushes its contents to disk. Lazy Writing is the process of updating the disk after the page has been changed in memory, so the application making the change to the file does not have to wait for the disk write to complete before proceeding. More than one page can be transferred on each write operation.
- The type is int.
- The unit is flushes/sec.

#### Cache Pin Reads/sec

- The frequency of reading data in to the cache preparatory to writing the data back to disk. Pages read in this fashion are pinned in memory at the completion of the read. While pinned, the physical address for a page in the cache is not altered.
- The type is int.
- The unit is reads/second.

#### Cache MDL Reads/sec

- The frequency of reads from cache pages that use a Memory Descriptor List (MDL) to access the data. The MDL contains the physical address of each page involved in the transfer, and thus can employ a hardware Direct Memory Access (DMA) device to effect the copy. The LAN Server uses this method for large transfers out of the server.
- The type is int.
- The unit is reads/second.

#### Cache Data Map Hits Dynamic Average

- A running average of the Data Map Hits % attribute.
- The type is int.
- The unit is percent.

#### Cache Data Flushes/sec

- The frequency with which the cache has flushed its contents to disk as the result of a request to flush or to satisfy a write-through file write request. More than one page can be transferred on each flush operation.
- The type is int.
- The unit is flushes/second.

#### Cache Async Copy Reads/sec

- The frequency of reads from cache pages that involve a memory copy of the data from the cache to the buffer for the application. The application regains control immediately even if the disk must be accessed to retrieve the page.
- The type is int.
- The unit is reads/second.

#### Cache MDL Read Hits Dynamic Average

- A running average of the Memory Descriptor List (MDL) Read Hits % attribute.



- The type is int.
- The unit is percent.

#### Cache Copy Read Hits %

- The percentage of cache copy read requests that hit the cache, that is, that did not require a disk read to provide access to the page in the cache. A copy read is a file read operation that is satisfied by a memory copy from a cache page to the buffer for the application. The LAN Redirector uses this method to retrieve cache information, as does the LAN Server for small transfers. This is a method used by the disk file systems as well.
- The type is int.
- The unit is percent.

#### Cache Sync Copy Reads/sec

- The frequency of reads from cache pages that involve a memory copy of the data from the cache to the buffer for the application. The file system does not regain control until the copy operation is complete, even if the disk must be accessed to retrieve the page.
- The type is int.
- The unit is reads/second.

#### Cache Async Data Maps/sec

- The frequency that an application uses a file system, such as NTFS or HPFS, to map a page of a file in to the cache to read the page, and does not wish to wait for the cache to retrieve the page if it is not in main memory.
- The type is int.
- The unit is maps/second.

#### Cache Data Map Pins/sec

- The frequency of data maps in the cache that resulted in pinning a page in main memory, an action usually preparatory to writing to the file on disk. While pinned, the physical address for a page in main memory and the virtual address in the cache is not altered.
- The type is int.
- The unit is pins/second.

#### Cache Sync Fast Reads/sec

- The frequency of reads from cache pages that bypass the installed file system and retrieve the data directly from the cache. Normally, file I/O requests invoke the appropriate file system to retrieve data from a file, but this path permits direct retrieval of cache data without file system involvement if the data is in the cache. Even if the data is not in the cache, one invocation of the file system is avoided. If the data is not in the cache, the request (application program call) waits until the data is retrieved from disk.
- The type is int.
- The unit is reads/second.

#### Cache Fast Read Not Possibles/sec

- The frequency of attempts by an Application Program Interface (API) function call to bypass the file system to get at cache data that could not be honored without invoking the file system.
- The type is int.
- The unit is attempts/second.

## Component: Active Tasks running on VCenter Server

Use the Index Service data set to monitor the creation of indices and the merging of indices by the indexing service. Indexing Service is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

### Dimensions

node

- The managed system name.
- The type is string.

Indexing Service Index (Unicode)

- A collection of all index information and stored properties for a particular group of file system directories.
- The type is string. This is a key dimension.

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

Indexing Service Index Name

- A collection of all index information and stored properties for a particular group of file system directories. Valid format is a text string of up to 64 characters.
- The type is string.

### Metrics

Indexing Service Deferred for Indexing

- Number of files not available and deferred for indexing.
- The type is int.
- The unit is files.

Indexing Service Total Number of Documents

- Total number of documents in the index.
- The type is int.
- The unit is documents.

Indexing Service Word Lists

- Number of word lists.
- The type is int.
- The unit is lists.

Indexing Service Indexing Speed MB/hr

- Speed of the indexing of file contents in MBs per hour.
- The type is int.
- The unit is megabytes/hour.

Indexing Service Running Queries

- Number of active query client connections.
- The type is int.

- The unit is count.

#### Indexing Service Merge Progress

- Percent merge complete for the current merge.
- The type is int.
- The unit is percent.

#### Indexing Service Binding Time mSec

- Average time spent binding to indexing filters.
- The type is int.
- The unit is milliseconds.

#### Indexing Service Index Size MB

- Size of the content index (ci files only) in MBs.
- The type is int.
- The unit is megabytes.

#### Indexing Service Total Number of Queries

- Total number of queries since the index was mounted. Note that this attribute is not available in IIS 4.0.
- The type is int.
- The unit is queries.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### Indexing Service Saved Indexes

- Number of saved indexes.
- The type is int.
- The unit is indexes.

#### Indexing Service Number of Documents Indexed

- Number of documents indexed since the current indexing session started.
- The type is int.
- The unit is documents.

#### Indexing Service Unique Keys

- Number of unique keys (words, etc) in the index.
- The type is int.
- The unit is keys.

#### Indexing Service Total Indexing Speed MB/hr

- Speed of indexing file contents and properties in MBs per hour.

- The type is int.
- The unit is megabytes/hour.

#### Indexing Service Files to be Indexed

- Number of files to be filtered and added to the index.
- The type is int.
- The unit is files.

### **Component: Pweformance Status Object**

The Performance Object Status attribute group contains information that reflects the status of other attribute groups so you can see the status of all of the performance objects that make up this application all at once. Each of these other performance attribute groups is represented by a row in this table (or other type of view). The status for an attribute group reflects the result of the last attempt to collect data for that attribute group, which allows you to see whether the agent is performing correctly. Unlike other attribute groups, the Performance Object Status attribute group does not reflect the state of the monitored application. This attribute group is most often used to determine why data is not available for one of the performance attribute groups.

#### **Dimensions**

##### Performance Status Object Last Collection Start

- The most recent time a data collection of this group started.
- The type is timestamp.

##### Performance Status Object Last Collection Finished

- The most recent time a data collection of this group finished.
- The type is timestamp.

##### Performance Status Object Object Name

- The name of the performance object.
- The type is string.

##### Performance Status Object Object Type

- The type of the performance object.
- The type is int.

##### Performance Status Object Object Status

- The status of the performance object.
- The type is int.

##### Performance Status Object Query Name

- The name of the attribute group.
- The type is string. This is a key dimension.

##### Timestamp

- This is the local time when the data was collected.
- The type is timestamp.

##### Performance Status Object System Name

- This is the managed system name of the agent.

- The type is string. This is a key dimension.

#### Performance Status Object Error Code

- The error code that is associated with the query.
- The type is int.

### **Metrics**

#### Performance Status Object Number of Collections

- The number of times this group has been collected since agent start.
- The type is int.
- The unit is occurrences.

#### Performance Status Object Average Collection Duration

- The average duration of all data collections of this group in seconds.
- The type is double.
- The unit is seconds.

#### Performance Status Object Cache Misses

- The number of times an external data request for this group was not available in the cache.
- The type is int.
- The unit is misses.

#### Performance Status Object Intervals Skipped

- The number of times a background data collection for this group was skipped because the previous collection was still running when the next one was due to start.
- The type is int.
- The unit is intervals.

#### Performance Status Object Refresh Interval

- The interval at which this group is refreshed in seconds.
- The type is int.
- The unit is seconds.

#### Performance Status Object Cache Hit Percent

- The percentage of external data requests for this group that were satisfied from the cache.
- The type is double.
- The unit is percent.

#### Performance Status Object Last Collection Duration

- The duration of the most recently completed data collection of this group in seconds.
- The type is double.
- The unit is seconds.

#### Performance Status Object Cache Hits

- The number of times an external data request for this group was satisfied from the cache.
- The type is int.
- The unit is hits.

## Component: HTTP Content Index

Information about queries made to an HTTP (HyperText Transport Protocol) server, such as the number of active queries, the current requests queued, and the percentage of queries found in the query cache.

### Dimensions

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

HTTP Control Index Parameter

- A special one for the front end to use as a column header.
- The type is string.

HTTP Control Index Value

- A special one for the front end to use as a column header.
- The type is string.

node

- The managed system name.
- The type is string.

### Metrics

HTTP Control Index Running Queries

- The current number of running queries.
- The type is int.
- The unit is queries.

HTTP Control Index Completed Cache Queries

- The number of completed queries in cache.
- The type is int.
- The unit is queries.

HTTP Control Index Total Requests Rejected

- The total number of query requests rejected.
- The type is int.
- The unit is requests.

HTTP Control Index Total Queries

- The total number of queries since the server started up.
- The type is int.
- The unit is queries.

Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.

- The type is int.
- The unit is row.

#### HTTP Content Index % Cache Hits

- The percentage of queries found in the query cache.
- The type is int.
- The unit is percent.

#### HTTP Control Index % Cache Misses

- The percentage of queries not found in the query cache.
- The type is int.
- The unit is percent.

#### HTTP Control Index Queries Per Minute

- The number of queries per minute.
- The type is int.
- The unit is queries/minute.

#### HTTP Control Index Current Requests Queued

- The current number of query requests queued.
- The type is int.
- The unit is queries.

### **Component: DHCP server**

Information about the Dynamic Host Configuration Protocol (DHCP) messages sent and received by the server, the average amount of processing time spent by the server per message packet, and the number of message packets dropped because of internal delays at the server.

#### **Dimensions**

##### DHCP Parameter

- A special one for the front end to use as a column header.
- The type is string.

##### DHCP Value

- A special one for the front end to use as a column header.
- The type is string.

##### node

- The managed system name.
- The type is string.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### **Metrics**

##### DHCP Active Queue Length

- The number of packets in the processing queue of the DHCP server.

- The type is int.
- The unit is packets.

#### DHCP Rate of DHCP Requests

- Rate of DHCP Requests received by the DHCP server.
- The type is int.
- The unit is requests/second.

#### DHCP Milliseconds Per Packet Average

- The average time per packet taken by the DHCP server to send a response. Valid values are positive integers in the range 0 to 2147483647 and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions.
- The type is int.
- The unit is milliseconds.

#### DHCP Rate of DHCP Nacks

- Rate of DHCP Nacks (negative acknowledgments) sent by the DHCP server.
- The type is int.
- The unit is nacks/second.

#### DHCP Packets Expired/sec

- Rate at which packets get expired in the DHCP server message queue.
- The type is int.
- The unit is packets/second.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### DHCP Packets Received/sec

- Rate at which packets are received by the DHCP server.
- The type is int.
- The unit is packets/second.

#### DHCP Release Rate

- Rate of DHCP Releases received by the DHCP server.
- The type is int.
- The unit is releases/second.

#### DHCP Conflict Check Queue Length

- The number of packets in the DHCP server queue waiting on conflict detection (ping).
- The type is int.
- The unit is packets.

#### DHCP Acks Rate



- Rate of DHCP Acks (acknowledgments) sent by the DHCP server.
- The type is int.
- The unit is acks/second.

#### DHCP Rate of DHCP Informs

- Rate of DHCP Informs received by the DHCP server.
- The type is int.
- The unit is informs/second.

#### DHCP Rate of DHCP Declines

- Rate of DHCP Declines received by the DHCP server.
- The type is int.
- The unit is declines/second.

#### DHCP Rate of DHCP Discovers

- Rate of DHCP Discovers received by the DHCP server.
- The type is int.
- The unit is discovers/second.

#### DHCP Duplicates Dropped/sec

- Rate at which the DHCP server received duplicate packets.
- The type is int.
- The unit is packets/second.

#### DHCP Rate of DHCP Offers

- Rate of DHCP Offers sent out by the DHCP server.
- The type is int.
- The unit is offers/second.

### **Component: Active Tasks running on VCenter Server**

Use the Printer data set to create situations that monitor information about each printer that is attached to your server. Printer is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### Printer Comment

- The comment. For example, SUBMITTED is an example of a comment.
- The type is string.

##### Printer Share Name (Unicode)

- The share name of the printer.
- The type is string.

##### Printer Separator File (Unicode)

- The file that contains the job separator page in UTF8.
- The type is string.

##### Printer Separator File

- The file that contains the job separator page. For example, to JSEP indicates the file that contains the job separator page.
- The type is string.

#### Printer Until Time

- The end time of the printer operation.
- The type is string.

#### Printer Data Type

- The data type used to record print jobs.
- The type is string.

#### Printer Comment (Unicode)

- Comment in UTF8. Valid format is a text string of up to 388 bytes.
- The type is string.

#### Printer Priority

- The priority of the job. Valid values are positive integers in the range 0 to 2147483647 and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions.
- The type is int.

#### Printer Default Priority

- The default priority value assigned to each print job.
- The type is int.

#### Printer node

- The managed system name.
- The type is string.

#### Printer Port Name

- The port name that the printer is connected to. For example, to LDEV1 specifies the port name for the printer.
- The type is string.

#### Printer Share Name

- The share name of the printer. Valid format is a text string of up to 32 characters. For example, to AGHQ01 specifies the share name of the printer.
- The type is string.

#### Printer Name (Unicode)

- The name of the printer in UTF8.
- The type is string. This is a key dimension.

#### Printer Location

- The location where the printer resides. For example, to AGH specifies a location where the printer resides.
- The type is string.

#### Printer Location (Unicode)

- The location of the printer in UTF8.
- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### Printer Name

- The name of the printer. Valid format is a text string of up to 64 characters. For example, LPT1 is the name of the printer.
- The type is string.

#### Printer Parameters

- The parameters of the print processor. Valid format is a text string of up to 64 characters. For example, MARGIN indicates a parameter.
- The type is string.

#### Printer Status of Job

- The status of the job. Valid format is a text string of up to 20 characters. For example, PRINTING specifies the status of the print job.
- The type is string.

#### Printer Print Processor

- The print processor that must be used. For example, WINPRINT indicates the print processor used.
- The type is string.

#### Printer Start Time

- The start time of the printer operation.
- The type is string.

#### Printer Driver Name

- The print driver that is being used. For example, HP Laserjet III is an example of a print driver.
- The type is string.

### **Metrics**

#### Printer Average Pages Per Minute

- The average pages printed per minute of the printer. This attribute is the 64-bit version of Average Pages Per Minute.
- The type is double.
- The unit is pages/minute.

#### Printer Number of Jobs (Superseded)

- The number of jobs in the queue.
- The type is int.
- The unit is jobs.

#### Printer Number of Jobs

- The number of jobs in the queue. This attribute is the 64-bit version of Number of Jobs.
- The type is double.
- The unit is jobs.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### Printer Average Pages Per Minute (Superseded)

- The average pages printed per minute of the printer.
- The type is int.
- The unit is pages/minute.

### Component: FTP Server Statistics

Information about traffic and connection activity about the FTP (File Transfer Protocol) Server, such as the current connections, the bytes received per second, and the total non-anonymous users connected.

#### Dimensions

##### FTP Server Statistics Parameter

- A special one for the front end to use as a column header.
- The type is string.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### node

- The managed system name. The format should be *hostname : agent\_code* . Examples include spark:KNT or deux.raleigh.ibm.com:KNT.
- The type is string.

##### FTP Server Statistics Value

- A special one for the front end to use as a column header.
- The type is string.

#### Metrics

##### FTP Server Statistics Total Non Anonymous Users Since FTP Start

- The total number of non-anonymous users connected since the FTP server was started.
- The type is int.
- The unit is users.

##### FTP Server Statistics Current Non Anonymous Users

- The number of non-anonymous users currently connected to the FTP server.
- The type is int.

- The unit is users.

#### FTP Server Statistics Logon Attempts Since FTP Start

- The number of logon attempts since the FTP server was started.
- The type is int.
- The unit is attempts.

#### FTP Server Statistics Connection Attempts Since FTP Start

- The number of connection attempts since the FTP server was started.
- The type is int.
- The unit is attempts.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### FTP Server Statistics Current Anonymous Users

- The number of anonymous users currently connected to the FTP server.
- The type is int.
- The unit is users.

#### FTP Server Statistics Total Bytes/sec

- The total number of Kbytes flowing through the FTP server per second. This includes both incoming and outgoing bytes. This number is a good indicator of how heavily your FTP server is loaded.
- The type is int.
- The unit is bytes/second.

#### FTP Server Statistics Bytes Received/sec

- The number of bytes received per second by the FTP server.
- The type is int.
- The unit is bytes/second.

#### FTP Server Statistics Maximum Non Anonymous Users

- The maximum number of non-anonymous users simultaneously connected to the FTP server.
- The type is int.
- The unit is users.

#### FTP Server Statistics Total Files Semt/Recieved by FTP Server

- The total number of files sent and received by the FTP server.
- The type is int.
- The unit is files.

#### FTP Server Statistics Files Sent by FTP Server

- The number of files sent by the FTP server.

- The type is int.
- The unit is files.

#### FTP Server Statistics Total Anonymous Users Since FTP Start

- The total number of anonymous users connected since the FTP server was started.
- The type is int.
- The unit is users.

#### FTP Server Statistics FTP Server Files Received

- The number of files received by the FTP server.
- The type is int.
- The unit is files.

#### FTP Server Statistics Current Connections

- The current number of connections to the FTP server.
- The type is int.
- The unit is connections.

#### FTP Server Statistics Maximum Anonymous Users

- The maximum number of anonymous users simultaneously connected to the FTP server.
- The type is int.
- The unit is users.

#### FTP Server Statistics Bytes Sent/sec

- The number of data bytes sent per second by the FTP server. Note: the value 2147483647 indicates Value\_Exceeds\_Maximum and the value -2147483648 indicates Value\_Exceeds\_Minimum.
- The type is int.
- The unit is bytes/sec.

#### FTP Server Statistics Maximum Connections

- The maximum number of simultaneous connections to the FTP server.
- The type is int.
- The unit is connections.

### **Component: Active Tasks running on VCenter Server**

Use the IP Address data set to obtain IP (Internet Protocol) address information.

#### **Dimensions**

##### IP Address IP Address

- An IP address associated with the network interface.
- The type is string. This is a key dimension.

##### IP Address DNS Name

- The Domain Name Server entry associated with the IP network address.
- The type is string.

##### IP Address Address Type

- An indicator as to whether the IP address is version 4 or version 6. The value for IPv4\_IPv6 relates to the Automatic Tunneling Pseudo Interface for the Windows Server 2003 and Windows XP operating systems. These pseudo processes include, System, Idle, and \_Total.
- The type is int.

#### IP Address Network Interface Name

- The name of the network interface.
- The type is string. This is a key dimension.

#### IP Address MAC Address

- The MAC address of the network interface.
- The type is string.

#### node

- The managed system name.
- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### IP Address Network Interface Name (Long)

- Long name of the network interface.
- The type is string.

### **Component: Active Tasks running on VCenter Server**

Use the File Change data set to monitor changes to your file system and to request notification when resources change. File Change is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### File Change Change File Name

- Selects monitoring for file name changes.
- The type is string.

##### File Change Date Created

- The Date that the file was created.
- The type is string.

##### File Change Change Last Access

- Selects monitoring for Last Access date/time changes.
- The type is string.

##### File Change Change Create

- Selects monitoring for Create Date/Time changes.
- The type is string.

##### File Change Change Attributes

- Selects monitoring for attribute changes.

- The type is string.

#### File Change Watch Directory

- The name of the watched directory. Any changes to any files in the directory will be detected. For example, C:FILESYSUSAGE . A filter value is required for this attribute in situations using either the File Change data set or the File Trend data set.
- The type is string.

#### Monitor all Conditions

- This attribute combines different types of filter criteria, allowing you to monitor the most recent file changes. You can use this attribute to trigger monitoring for all of the following attributes: Change Attributes, Change Create, Change Directory Name, Change File Name, Change Last Access, Change Last Write, Change Security, and Change Size. You can also use the above-listed attributes in individual situations to monitor for one or more conditions, such as the latest change to the file name or the last time a file was accessed.
- The type is string.

#### File Change Watch Tree

- The entire Watch Tree or only the directory. y = to select monitoring for the entire watch tree. n = to select monitoring for the directory only. For example, y indicates that you are monitoring the entire watch tree.
- The type is string.

#### File Change Change Security

- Selects monitoring for security code changes.
- The type is string.

#### Modify Time

- The time that the file was last modified.
- The type is string.

#### File Change Date Time Created

- The date and time at which the file was created.
- The type is timestamp.

#### Modify Date

- The date that the file was last modified.
- The type is string.

#### File Change Watch Directory (Unicode)

- The name of the Watch Directory in UTF8. Any changes to any files in the directory will be detected. A filter value is required for this attribute in situations using either the File Change data set or the File Trend data set.
- The type is string.

#### File Change Watch File (Unicode)

- The path and name of the Watch File in UTF8. A filter value is required for this attribute in situations using either the File Change data set or the File Trend data set.
- The type is string. This is a key dimension.



#### File Change Change Size

- Selects monitoring for size changes. Valid value is a single character (y=Yes or n=No).
- The type is string.

#### File Change Attributes

- The current attributes for the file. When used in a situation, this value does not act as a trigger condition for the situation. This value can be used as a filter condition after the situation has already been triggered by one of the predefined monitor conditions.
- The type is int.

#### File Change Change Directory Name

- Selects monitoring for directory name changes.
- The type is string.

#### Date Time Last Modified

- The date and time the file was last modified. For example, 1101009130500000 expresses October 9, 2010, 1:05:00 pm.
- The type is timestamp.

#### File Change Watch File

- The path and name of the Watch File. For example, usage. log. A filter value is required for this attribute in situations using either the File Change data set or the File Trend data set.
- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### File Change Time Created

- The time that the file was created.
- The type is string.

#### File Change Action

- The type of change that occurred most recently to the directory or to the file. The numbers 1 through 5 determine the type of change that recently occurred.
- The type is int.

#### File Change Change Last Write

- Selects monitoring for Last Write date/time changes.
- The type is string.

#### File Change Boundry Alignment

- This is for boundary alignment when multiple occurrences of the row are allocated.  
\*IMPORTANT\* if additional attributes are added or existing attributes change in size, this attribute may need to be adjusted.
- The type is string.

node

- The managed system name. The format should be *hostname : agent\_code* . Examples include spark:KNT or deux.raleigh.ibm.com:KNT.
- The type is string.

## Metrics

### File Trend Size Change Last Interval

- The average absolute file size change over the last interval. This attribute is the 64-bit version of Size\_Change\_LastInterval. Note: -1 indicates Unknown and 9223372036854775807 indicates Value\_Exceeds\_Maximum.
- The type is double.
- The unit is bytes.

### File Trend Size Change Last Hour (Superseded)

- The average absolute file size change over the last hour. Note: -1 indicates Unknown and 2147483647 indicates Value\_Exceeds\_Maximum.
- The type is int.
- The unit is bytes.

### Sampling Interval

- The time between successive samples. Note: -1 indicates Unknown and 2147483647 indicates Value\_Exceeds\_Maximum.
- The type is int.
- The unit is seconds.

### File Trend Size Change Total (Superseded)

- The absolute file size change since monitoring began expressed as a percentage. Valid values are positive integers in the range of 0 to 2147483647 and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions. Note: -1 indicates Unknown and 2147483647 indicates Value\_Exceeds\_Maximum.
- The type is int.
- The unit is percent.

### File Trend Size Change Total

- The absolute file size change since monitoring began expressed as a percentage. This attribute is the 64-bit version of Size Change Total. Note: -1 indicates Unknown and 9223372036854775807 indicates Value\_Exceeds\_Maximum.
- The type is double.
- The unit is percent.

### File Trend % Change Last Hour

- The percentage rate of growth over the last hour. Note: -1 indicates Unknown and 2147483647 indicates Value\_Exceeds\_Maximum.
- The type is int.
- The unit is percent.

### File Trend Percent Used

- The percentage of the entire disk resource that a particular file takes. Note: -1 indicates Unknown and 2147483647 indicates Value\_Exceeds\_Maximum.
- The type is int.

- The unit is percent.

#### File Trend % Change Last Interval

- The percentage rate of growth over the last interval. Note: -1 indicates Unknown and 2147483647 indicates Value\_Exceeds\_Maximum.
- The type is int.
- The unit is percent.

#### File Trend % Change Total

- The percentage rate of growth change since monitoring began. Note: -1 indicates Unknown and 2147483647 indicates Value\_Exceeds\_Maximum.
- The type is int.
- The unit is percent.

#### Total Hits (Superseded)

- The total number of file/directory changes since monitoring began. Note: the value 2147483647 indicates Value\_Exceeds\_Maximum and the value -2147483648 indicates Value\_Exceeds\_Minimum.
- The type is int.
- The unit is hits.

#### File Change Total Hits

- The total number of file/directory changes since monitoring began. This attribute is the 64-bit version of Total\_Hits. Note: the value 9223372036854775807 indicates Value\_Exceeds\_Maximum and the value -9223372036854775808 indicates Value\_Exceeds\_Minimum.
- The type is double.
- The unit is hits.

#### Sampling Number

- The number of intervals that are sampled to get an average.
- The type is int.
- The unit is number.

#### File Trend % Change Average

- The percentage of change for the averaging take sample interval. Note: -1 indicates Unknown and 2147483647 indicates Value\_Exceeds\_Maximum.
- The type is int.
- The unit is percent.

#### File Trend Size Change Average

- The average absolute file size change. This attribute is the 64-bit version of Size Change Average. Note: -1 indicates Unknown and 9223372036854775807 indicates Value\_Exceeds\_Maximum.
- The type is double.
- The unit is bytes.

#### File Trend Free Space Exhausted Hours

- The time until current free space on the volume is exhausted based on the size change rate over the last hour. If there is no size change detected over the last hour then the Free Space Exhausted Hours value is zero(0). Note: -1 indicates Unknown and 2147483647 indicates Value\_Exceeds\_Maximum.
- The type is int.
- The unit is hours.

#### File Trend Size Change Last Hour

- The average absolute file size change over the last hour. This attribute is the 64-bit version of Size Change LastHour. Note: -1 indicates Unknown and 9223372036854775807 indicates Value\_Exceeds\_Maximum.
- The type is double.
- The unit is bytes.

#### File Trend Size Change Last Interval (Superseded)

- The average absolute file size change over the last interval. Valid values are positive integers in the range 0 to 2147483647 and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions. Note: -1 indicates Unknown and 2147483647 indicates Value\_Exceeds\_Maximum.
- The type is int.
- The unit is bytes.

#### Current Size (Superseded)

- The current size of the file. When used in a situation, this value does not act as a trigger condition for the situation. This value can be used as a filter condition after the situation has already been triggered by one of the predefined monitor conditions. Note: the value 2147483647 indicates Value\_Exceeds\_Maximum and the value -2147483648 indicates Value\_Exceeds\_Minimum.
- The type is int.
- The unit is bytes.

#### File Change Current Size

- The current size of the file. When used in a situation, this value does not act as a trigger condition for the situation. This value can be used as a filter condition after the situation has already been triggered by one of the predefined monitor conditions. This attribute is the 64-bit version of Current\_Size. Note: the value 9223372036854775807 indicates Value\_Exceeds\_Maximum and the value -9223372036854775808 indicates Value\_Exceeds\_Minimum.
- The type is double.
- The unit is bytes.

#### Row Number

- Row Number.
- The type is int.
- The unit is row.

#### File Trend Size Change Average (Superseded)

- The average absolute file size change. Note: -1 indicates Unknown and 2147483647 indicates Value\_Exceeds\_Maximum.
- The type is int.
- The unit is bytes.

## Component: TCP Statistics

Information about monitoring connection statistics and segment traffic for data using the TCP protocol. TCP Statistics is a single-instance data set. This data set reports IPv4 statistics. IPv6 statistics are reported separately on Windows 2003.

### Dimensions

node

- The managed system name.
- The type is string.

TCP Statistics Parameter

- A special one for the front end to use as a column header.
- The type is string.

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

TCP Statistics Value

- A special one for the front end to use as a column header.
- The type is string.

### Metrics

TCP Statistics Segments Received/sec

- The rate that segments are received, including those received in error. This count includes segments received on currently established connections.
- The type is int.
- The unit is segments/second.

TCP Statistics Segments Sent/sec

- The rate that segments are sent, including those on current connections. This count excludes those containing only retransmitted bytes.
- The type is int.
- The unit is segments/second.

TCP Statistics Segments Retransmitted/sec

- The rate that segments are retransmitted, that is, segments transmitted containing one or more previously transmitted bytes.
- The type is int.
- The unit is segments/second.

TCP Statistics Connections Reset

- The number of times TCP connections have made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state. Valid format is a text string of up to 64 characters.
- The type is int.
- The unit is connections.

#### TCP Statistics Connection Failures

- The number of times TCP connections have made a direct transition to the CLOSED state from the SYN-SENT state or from the SYN-RCVD state. This number also includes the number of times TCP connections have made a direct transition to the LISTEN state from the SYN-RCVD state.
- The type is int.
- The unit is connections.

#### TCP Statistics Connections Established

- The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.
- The type is int.
- The unit is connections.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### TCP Statistics Connections Active

- The number of times TCP connections have made a direct transition to the SYN-SENT state from the CLOSED state.
- The type is int.
- The unit is connections.

#### TCP Statistics Connections Passive

- The number of times TCP connections have made a direct transition to the SYN-RCVD state from the LISTEN state.
- The type is int.
- The unit is connections.

#### TCP Statistics Segments/sec

- The rate that TCP segments are sent or received using the TCP protocol.
- The type is int.
- The unit is segments/second.

### **Component: Processor Summary**

Information about high and low processor information for one server.

#### **Dimensions**

##### Processor Summary High Process ID

- The process ID of the process with the highest processor utilization on this server.
- The type is int.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.

- The type is timestamp.

#### Processor Summary Lowest Processor Utilizer

- Name of the processor with the lowest utilization.
- The type is string.

#### Processor Summary High Processor Name

- Name of the processor with the highest utilization.
- The type is string.

#### node

- The managed system name. The format should be *hostname : agent\_code* . Examples include spark:KNT or deux.raleigh.ibm.com:KNT.
- The type is string.

#### Processor Summary High Process Name

- The name of the process with the highest processor utilization on this server.
- The type is string.

### Metrics

#### PProcessor Summary rocessor Utilization Difference

- On a multi-processor computer, the difference in processor total utilization between the processor with the highest total utilization percentage and the processor with the lowest total utilization percentage. On a single processor computer this value is 0.
- The type is int.
- The unit is percent.

#### Processor Summary Percent Interrupt Time for High Processor

- Percent of interrupt time for the High Processor. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Processor Summary Inturrupts/Sec for High Processor

- Number of interrupts per second for the High Processor. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is interrupts/sec.

#### Processor Summary Percent User Time for High Processor

- Percent of user time for the High Processor. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Processor Summary Interrupts/sec for Low Processor

- Number of interrupts per second for the Low Processor.
- The type is int.

- The unit is interrupts/second.

#### Processor Summary High Process Utilization

- The total processor utilization of the process with the highest processor utilization on this server.
- The type is int.
- The unit is percent.

#### Processor Summary High Process Average Utilization

- The average processor utilization of the process with the highest processor utilization on this server.
- The type is int.
- The unit is percent.

#### Processor Summary Processor Interrupt Difference

- On a multi-processor computer, the difference in processor utilization for interrupt handling between the processor with the highest total utilization percentage and the processor with the lowest total utilization percentage.
- The type is int.
- The unit is percent.

#### Processor Summary Processor User Difference

- On a multi-processor computer, the difference in processor utilization for user handling between the processor with the highest total utilization percentage and the processor with the lowest total utilization percentage. Note that the processor with the lower total utilization has a user processing time higher than the higher processor. When this occurs, this attribute has a negative value. On a single processor computer this value is 0.
- The type is int.
- The unit is percent.

#### Processor Summary Percent Priviledged Time for High Processor

- Percent of privileged time for the High Processor. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Processor Summary Percent Interrupt Time for Low Processor

- Percent of interrupt time for the Low Processor. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Processor Summary Processor Privileged Difference

- On a multi-processor computer, the difference in processor utilization for privileged handling between the processor with the highest total utilization percentage and the processor with the lowest total utilization percentage. Note that the processor with the lower total utilization might have a privileged processing time higher than the higher processor. When this occurs, this attribute has a negative value. On a single processor computer this value is 0.
- The type is int.



- The unit is percent.

#### Processor Summary Percent Privileged Time for Low Processor

- Percent of privileged time for the Low Processor. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Processor Summary Percent User Time for Low Processor

- Percent of user time for the Low Processor. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Processor Summary Low Processor Percent Utilization

- Percent of total processor time for the Low Processor. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Processor Summary Percent Total Time for High Processor

- Percent of total processor time for the High Processor. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

### **Component: Active Tasks running on VCenter Server**

Use the Print Job data set to create situations that monitor information about each print job owned by a specific printer that is attached to your server. Print Job is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### Print Job Machine Name

- The machine that created the job. For example, AGHQ01 is an example of a machine name.
- The type is string.

##### Print Job Notify Name

- The user to notify about the job. For example, MBROWN is an example of a user name.
- The type is string.

##### Print Job Printer Name

- The name of the printer. Valid format is a text string of up to 64 characters. For example, to LPT1 specifies the name of the printer.
- The type is string.

##### Print Job node

- The managed system name. The format should be *hostname : agent\_code* . Examples include spark:KNT or deux.raleigh.ibm.com:KNT.
- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### Print Job Data Type

- The data type used to record print jobs. For example, tEXT is an example of a data type.
- The type is string.

#### Print Job Job Status

- The status of the job. Valid format is a text string of up to 20 characters. For example, to PRINTING indicates the status of the job.
- The type is string.

#### Print Job Priority

- The priority of the job. Valid values are positive integers in the range 0 to 2147483647 and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions.
- The type is int.

#### Print Job Printer Name (Unicode)

- The name of the printer in UTF8.
- The type is string. This is a key dimension.

#### Print Job Position

- The position of the job in the print queue.
- The type is int.

#### Print Job Parameters

- The parameters of the print processor. Valid format is a text string of up to 64 characters. For example, MARGIN is an example of a parameter setting.
- The type is string.

#### Print Job User Name

- The user who owns the job. Valid format is a text string of up to 32 characters. For example, SMITH is an example of a user name.
- The type is string.

#### Print Job Print Processor

- The print processor that must be used. For example, WINPRINT is an example of a print processor.
- The type is string.

#### Print Job Driver Name

- The print driver that is being used. For example, to HP Laserjet III is an example of a print driver name.
- The type is string.

#### Print Job Time Submitted

- The time when the job was submitted.

- The type is timestamp.

Print Job User Name (Unicode)

- The user who owns the job in UTF8.
- The type is string.

Print Job Notify Name (Unicode)

- The user to notify about the job in UTF8.
- The type is string.

Print Job Document Name

- The name of the document in print. For example, KNTDOC is an example of a document name.
- The type is string.

Print Job Document Name (Unicode)

- The name of the document in print in UTF8.
- The type is string. This is a key dimension.

## **Metrics**

Print Job Size (Superseded)

- The size of the print job.
- The type is int.
- The unit is requests.

Print Job Pages Printed (Superseded)

- The number of pages that have printed.
- The type is int.
- The unit is pages.

Print Job Size

- The size of the print job. This attribute is the 64-bit version of Size.
- The type is double.
- The unit is requests.

Print Job Total Pages (Superseded)

- The number of pages of the job.
- The type is int.
- The unit is pages.

Print Job Time Elapsed (Superseded)

- The time elapsed, in seconds, that has elapsed since the job began printing.
- The type is int.
- The unit is seconds.

Print Job Time Elapsed

- The time elapsed, in seconds, that has elapsed since the job began printing. This attribute is the 64-bit version of Time Elapsed.

- The type is double.
- The unit is seconds.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### Print Job Total Pages

- The number of pages of the job. This attribute is the 64-bit version of Total Pages.
- The type is double.
- The unit is pages.

#### Print Job Pages Printed

- The number of pages that have printed. This attribute is the 64-bit version of Pages Printed.
- The type is double.
- The unit is pages.

### **Component: Active Tasks running on VCenter Server**

Use the Processor data set to create situations that monitor information about each processor on the computer. Processor is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### Processor Socket Designation

- Type of chip socket used on the circuit card.
- The type is string.

##### Processor L2 Cache Size

- L2 cache size of the processor. This value is presented in KB.
- The type is string.

##### Processor Current Clock Speed

- Current speed of the processor, in MHz.
- The type is string.

##### Processor Manufacturer

- Name of the processor manufacturer.
- The type is string.

##### Processor Data Width

- Processor data width, in bits.

- The type is string.

#### Processor Load Percentage

- Load capacity of each processor, averaged to the last second. Processor loading refers to the total computing burden for each processor at one time.
- The type is string.

#### Processor Description

- Processor description.
- The type is string.

#### node

- The managed system name.
- The type is string.

#### Processor Power Management Support

- If TRUE, the power of the device can be managed, which means that it can be put in to suspend mode, and so on. The property does not indicate that power management features are enabled, but it does indicate that the logical device power can be managed.
- The type is string.

#### Processor ID

- Processor information that describes the processor features. For an x86 class CPU, the field format depends on the processor support of the CPUID instruction. If the instruction is supported, the property contains 2 (two) DWORD formatted values. The first is an offset of 08h-0Bh, which is the EAX value that a CPUID instruction returns with input EAX set to 1. The second is an offset of 0Ch-0Fh, which is the EDX value that the instruction returns. Only the first two bytes of the property are significant and contain the contents of the DX register at CPU reset all others are set to 0 (zero), and the contents are in DWORD format.
- The type is string.

#### Processor Device ID

- Unique identifier of a processor on the system.
- The type is string. This is a key dimension.

#### Processor Address Width

- Processor address width, in bits. This represents the size of a pointer type on the processor. On a 32-bit processor, the value is 32 and on a 64-bit processor it is 64.
- The type is string.

#### Processor Name

- Processor name.
- The type is string.

#### node

- The managed system name.
- The type is string.

#### Processor Version

- Processor revision number that depends on the architecture.

- The type is string.

#### Processor Maximum Clock Speed

- Maximum speed of the processor, in MHz.
- The type is string.

#### Processor Instance Name

- The processor instance name. Valid format is a text string of up to 64 characters. For example, SYS1.
- The type is string. This is a key dimension.

### Metrics

#### Processor Percent Time Deferred Procedure Calls

- The percentage of processor time spent processing Deferred Procedure Calls (DPCs) during the sample interval. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Row Number

- Row Number.
- The type is int.
- The unit is row.

#### Processor Average Interrupts Processor per Sec

- The average number of interrupts a processor has processed per second.
- The type is int.
- The unit is interrupts/second.

#### Processor Deferred Procedure Calls Avoidance

- The rate at which Deferred Procedure Calls (DPCs) on all processors were avoided. (DPCs are interrupts that run at a lower priority than standard interrupts). This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
- The type is int.
- The unit is bypasses/second.

#### Processor Percent processor busy instructions non-idle privilege mode

- The percentage of elapsed time that a processor has been busy executing instructions in non-idle privileged mode. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Processor Percent time Processor Instructions Non Idle Mode

- The percentage of elapsed time a processor has been busy executing instructions in non-idle user mode. The percent value might exceed 100 on multiple processor systems. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.

- The unit is percent.

#### Percent Processor Time Hardware Interrupts

- The percentage of processor time spent processing hardware interrupts during the sample interval. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Processor Deferred Procedure Calls Added Rate

- The rate at which Deferred Procedure Calls (DPCs) are added to the DPC queue for the processor between the timer ticks of the processor clock. This is not the number of DPCs in the queue.
- The type is int.
- The unit is calls.

#### Processor Deferred Procedure Call Added Rate

- The overall rate at which Deferred Procedure Calls (DPCs) are added to the DPC queue for the processor. This is not the number of DPCs in the queue. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
- The type is int.
- The unit is calls/second.

#### Processor Percent Time Processor Busy Non Idle Threads

- The percentage of elapsed time that a processor has been busy executing non-idle threads. Valid values are positive integers in the range 0 to 100 (expressing a percentage) and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions. The percent value might exceed 100 on multiple processor systems. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Processor Kernel APC Interruption Avoidance

- The rate at which Kernel APC interrupts were avoided. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
- The type is int.
- The unit is bypasses/second.

### **Component: Windows Custom Scripts**

Statistic data collected using custom scripts. It contains one row for each defined script.

#### **Dimensions**

##### Custom Scripts Execution Start

- The time when the last execution of this script started.
- The type is timestamp.

##### Custom Scripts Runtime Property File Name

- The name of the property file.

- The type is string. This is a key dimension.

#### Custom Scripts Runtime Attribute Name

- The attribute name that is defined in the properties file. The attribute is used for metric identification.
- The type is string. This is a key dimension.

#### Custom Scripts Runtime Custom Name

- The custom name that is defined in the properties file. It is used for custom reporting.
- The type is string.

#### Custom Scripts Script Standard Error

- Script Standard Error in a unique row.
- The type is string.

#### Custom Scripts Runtime Return Code

- Integer value returned by the Script.
- The type is int.

#### Custom Scripts Runtime Timestamp

- This is the local time when the data was collected.
- The type is timestamp.

#### Custom Scripts Status Code

- The status of the script. It includes general errors, configuration errors, the status or the execution code returned by the Factory Script Provider.
- The type is int.

#### Custom Scripts Runtime Script Path Name

- The fully qualified path of the script.
- The type is string.

#### Custom Scripts Runtime String Standard Output

- Script Standard Output in String Format.
- The type is string.

#### Custom Scripts Property Group

- The name of the property group.
- The type is string.

#### Custom Scripts Custom Label Float #5

- Label for custom floating point attribute #5.
- The type is string.

#### Custom Scripts Custom Label Float #3

- Label for custom floating point attribute #3.
- The type is string.



#### Custom Label Float #4

- Label for custom floating point attribute #4.
- The type is string.

#### Custom Scripts Runtime Custom attribute String #5

- Custom string attribute #5.
- The type is string.

#### Custom Scripts Script Name

- The name of the script.
- The type is string.

#### Output Type

- Standard output type of the script.
- The type is int.

#### Custom Scripts Runtime Custom attribute String #1

- Custom string attribute #1.
- The type is string.

#### Custom Label String #1

- Label for custom string attribute #1.
- The type is string.

#### Custom Scripts Runtime Custom attribute String #2

- Custom string attribute #2.
- The type is string.

#### Custom Scripts Runtime System Name

- This is the managed system name of the agent.
- The type is string.

#### Custom Scripts Custom Label String #2

- Label for custom string attribute #2.
- The type is string.

#### Custom attribute String #3

- Custom string attribute #3.
- The type is string.

#### Custom Scripts Custom Label String #3

- Label for custom string attribute #3.
- The type is string.

#### Custom Scripts Runtime Custom attribute String #4

- Custom string attribute #4.
- The type is string.

#### Custom Label String #4

- Label for custom string attribute #4.
- The type is string.

#### Custom Scripts Custom Label Float #1

- Label for custom floating point attribute #1.
- The type is string.

#### Custom Scripts Custom Label String #5

- Label for custom string attribute #5.
- The type is string.

#### Custom Scripts Custom Label Float #2

- Label for custom floating point attribute #2.
- The type is string.

#### Custom Scripts Runtime Row Number

- Output row number.
- The type is int. This is a key dimension.

#### Custom Scripts Custom Label Integer #2

- Label for custom integer attribute #2.
- The type is string.

#### Custom Scripts Custom Label Integer #3

- Label for custom integer attribute #3.
- The type is string.

#### Custom Scripts Custom Label Integer #4

- Label for custom integer attribute #4.
- The type is string.

#### Custom Scripts Custom Label Integer #5

- Label for custom integer attribute #5.
- The type is string.

#### Custom Scripts Custom Label Integer #1

- Label for custom integer attribute #1.
- The type is string.

### **Metrics**

#### Custom Scripts Intervals Skipped

- The count of occurrences where an execution of this script is skipped because the previous execution is still running.
- The type is int.
- The unit is intervals.

#### Custom Scripts Execution Duration

- The duration of the last execution of this script, in seconds. When timing out, the value of the configured timeout is returned.
- The type is int.
- The unit is seconds.

#### Custom Scripts Runtime Custom attribute Integer #3

- Custom integer attribute #3.
- The type is double.
- The unit is value.

#### Custom Scripts Runtime Custom attribute Float #2

- Custom floating point (2 decimals) attribute #2.
- The type is double.
- The unit is value.

#### Custom Scripts Runtime Custom attribute Integer #4

- Custom integer attribute #4.
- The type is double.
- The unit is value.

#### Custom Scripts Runtime Custom attribute Float #3

- Custom floating point (2 decimals) attribute #3.
- The type is double.
- The unit is value.

#### Custom Scripts Refresh Interval

- The interval, in seconds, that the agent attempts to start this script.
- The type is int.
- The unit is seconds.

#### Custom Scripts Runtime Custom attribute Integer #5

- Custom integer attribute #5.
- The type is double.
- The unit is value.

#### Custom Scripts Runtime Custom attribute Float #1

- Custom floating point (2 decimals) attribute #1.
- The type is double.
- The unit is value.

#### Custom Scripts Runtime Custom attribute Float #4

- Custom floating point (2 decimals) attribute #4.
- The type is double.
- The unit is value.

#### Custom Scripts Runtime Custom attribute Float #5

- Custom floating point (2 decimals) attribute #5.

- The type is double.
- The unit is value.

#### Float Standard Output

- Script Output in Floating Point Format (2 decimals).
- The type is double.
- The unit is value.

#### Custom Scripts Runtime Integer Standard Output

- Script Output in Integer Format.
- The type is double.
- The unit is value.

#### Custom Scripts Runtime Custom attribute Integer #1

- Custom integer attribute #1.
- The type is double.
- The unit is value.

#### Custom Scripts Average Execution Duration

- The average duration, in seconds, of all the execution of the script.
- The type is double.
- The unit is seconds.

#### Custom Scripts Number of Executions

- The count of execution attempts of this script since agent started.
- The type is int.
- The unit is occurrences.

#### Custom Scripts Runtime Custom attribute Integer #2

- Custom integer attribute #2.
- The type is double.
- The unit is value.

### **Component: Active Tasks running on VCenter Server**

Use the Paging File data set to the page files of the system. Paging File is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### Page File Name

- The name of a page file. Valid format is a text string of up to 64 characters. For example, PAGING.
- The type is string.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

node

- The managed system name.
- The type is string.

#### Pagefile Name (Unicode)

- The instance name in UTF8.
- The type is string. This is a key dimension.

### Metrics

#### Paging File Percent Used

- The amount of a paging file in use. Valid values are positive integers in the range 0 to 100 (expressing a percentage) and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions.
- The type is int.
- The unit is percent.

#### Peak Page File Usage

- The peak amount of the Page File instance used in percent.
- The type is int.
- The unit is percent.

#### Row Number

- Row Number.
- The type is int.
- The unit is row.

### Component: Processor Count Information

Information about the total count information for all processors on a system. This information represents overall system activity.

#### Dimensions

##### System Number of Logical Processors

- The number of the logical processors on the system.
- The type is int.

##### System Operating System Type

- The marketed operating system name of the installed OS version. 2, Windows\_2000 for version 5. 0, and Windows\_XP for version 5. 1.
- The type is string.

##### System Network Address

- The host address of a system.
- The type is string.

##### System Parameter

- A special one for the front end to use as a column header.
- The type is string.

##### System Hyper-Threading

- Indicates whether or not the Hyper-Threading technology is active on the computer's processors.
- The type is int.

#### System Number of Processors

- The total number of processor cores in a system.
- The type is int.

#### System Value

- A special one for the front end to use as a column header.
- The type is string.

#### System Processor Type

- The processor type of a system.
- The type is int.

#### node

- The managed system name.
- The type is string.

#### System Operating System Version

- The operating system version number of a system.
- The type is string.

#### System User Name (Unicode)

- The name of a unique user account in UTF8.
- The type is string.

#### System User Name

- The name of a unique user account.
- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### System Total Memory Size (MB)

- The number of MBs of installed random access memory (RAM) in the computer.
- The type is int.

#### System Network Address IPv6

- The IPv6 host address for the computer. Note that the value No\_DNS\_Entry is a valid value.
- The type is string.

#### System Page Size (Bytes)

- The size of a page of virtual memory on a system in bytes.
- The type is int.

## Metrics

### System Up Time (Seconds) (Superseded)

- The total amount of time the system has been operational since it was last started.
- The type is int.
- The unit is seconds.

### System Processor Queue Length (Threads)

- The total number of threads waiting for processor time on a system.
- The type is int.
- The unit is threads.

### System Calls/sec

- The number of calls made to system service routines per second.
- The type is int.
- The unit is calls/second.

### System % Total User Time

- The average percentage of time all processors have been busy executing instructions in user mode. Note that the attribute value is averaged in synch with the situation or historical collection interval.  
The % Total User Time is the average percentage of time spent in User mode by all processors. On a multiprocessor system, if all processors are always in User mode this is 100%, if all processors are 50% in User mode this is 30% and if one-fourth of the processors are in User mode this is 25%. Applications execute in User Mode, as do subsystems like the window manager and the graphics engine. Code executing in User Mode cannot damage the integrity of the Windows NT Executive, Kernel, and device drivers. Unlike some early operating systems, Windows NT uses process boundaries for subsystem protection in addition to the traditional protection of User and Privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Windows NT on behalf of an application may appear in other subsystem processes in addition to the Privileged Time in the application process.
- The type is int.
- The unit is percent.

### System File Read Bytes/sec (Superseded)

- The average number of bytes the system transferred per second for all file system read operations.
- The type is int.
- The unit is bytes/second.

### System Exception Dispatches/sec

- The rate of exceptions dispatched by the system. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
- The type is int.
- The unit is dispatches/second.

### System File Write Bytes/sec

- The aggregate of the bytes transferred for all the file system write operations on the computer. This attribute is the 64-bit version of File Write Bytes/sec.
- The type is double.

- The unit is bytes/second.

#### System File Write Operations/sec

- The total number of file write operations the system has executed per second.
- The type is int.
- The unit is operations/second.

#### System File Write Bytes/sec (Superseded)

- The total number of bytes the system has transferred per second for all file write operations.
- The type is int.
- The unit is bytes/second.

#### System Processor Queue Length Excess

- The number of processor queue length requests in excess of the number of processors in the system. This indicates that the processor(s) are not able to service the work load that is requested of them. Calculated as Processor Queue Length - number of processes.
- The type is int.
- The unit is requests.

#### Row Number

- Row Number.
- The type is int.
- The unit is row.

#### System File Data Operations/sec

- The total number of read and write operations the system has executed per second.
- The type is int.
- The unit is operations/second.

#### System File Read Operations/sec

- The total number of file read operations the system has executed per second.
- The type is int.
- The unit is operations/second.

#### System File Control Bytes/sec

- The aggregate of bytes transferred for all file system operations that are neither reads nor writes. These operations usually include file system control requests or requests for information about device characteristics or status. This attribute is the 64-bit version of File Control\_Bytes/sec.
- The type is double.
- The unit is bytes/second.

#### System % Total Processor Time

- The % Total Processor Time is the average percentage of time that all the processors on the system are busy executing non-idle threads. On a multiprocessor system, if all processors are always busy this is 100%, if all processors are 50% busy this is 50% and if one-fourth of the processors are busy this is 25%. It can be viewed as the fraction of the time spent doing useful work. Each processor is assigned an Idle thread in the Idle process which consumes those



unproductive processor cycles not used by any other threads. Note: -1 indicates Unknown. Note that the attribute value is averaged in synch with the situation or historical collection interval.

- The type is int.
- The unit is percent.

#### System Up Time (Seconds)

- The total time (in seconds) that the computer has been operational since it was last started. This attribute is the 64-bit version of System Up Time (Seconds).
- The type is double.
- The unit is seconds.

#### System Alignment Fixups/sec

- The rate of alignment faults fixed by the system. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
- The type is int.
- The unit is fixups/second.

#### System Up Time (Days)

- Total Time (in days) that the computer has been operational since it was last started. Valid values include `Less_Than_One_Day`.
- The type is int.
- The unit is days.

#### System Page File Size (MB)

- Size of the system paging file, in MB.
- The type is int.
- The unit is megabytes.

#### System File Read Bytes/sec

- The aggregate of the bytes transferred for all the file system read operations on the computer. This attribute is the 64-bit version of File Read Bytes/sec.
- The type is double.
- The unit is bytes/second.

#### System Floating Emulations/sec

- The rate of floating emulations performed by the system. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
- The type is int.
- The unit is emulations/second.

#### System Total Interrupts/sec

- The rate at which the system is receiving and servicing hardware interrupts.
- The type is int.
- The unit is interrupts/second.

#### System Context Switches/sec

- The total number of context switches that have occurred on a system per second.

- The type is int.
- The unit is switches/second.

#### System % Total Privileged Time

- The total percentage of elapsed time a system has been busy executing instructions in privileged mode. Note: -1 indicates Unknown. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### System File Control Operations/sec

- The total number of file control operations the system has executed per second. This includes operations which are neither reads nor writes, such as file system control requests.
- The type is int.
- The unit is operations/second.

#### System File Control Bytes/sec (Superseded)

- The total number of bytes the system has transferred per second for all file control operations.
- The type is int.
- The unit is bytes/second.

### **Component: IP Statistics**

Information about traffic and fragmentation statistics for data using the IP (Internet Protocol) protocol.

#### **Dimensions**

node

- The managed system name.
- The type is string.

IP Statistics Parameter

- A special one for the front end to use as a column header.
- The type is string.

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

IP Statistics Column Header

- A special one for the front end to use as a column header.
- The type is string.

#### **Metrics**

IP Statistics Fragments Received/sec

- The rate that IP fragments that need to be re-assembled at this entity are received.
- The type is int.
- The unit is fragments/second.

#### IP Statistics Fragmentation Failures

- The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be. This happens due to setting the Don't Fragment flag.
- The type is int.
- The unit is datagrams.

#### IP Statistics Datagrams Received Delivered/sec

- The rate that input datagrams are successfully delivered to IP user-protocols (including ICMP).
- The type is int.
- The unit is datagrams/second.

#### IP Statistics Datagrams Received Address Errors

- The number of input datagrams discarded because the IP address in their IP header destination field was not a valid address to be received at this entity. This count includes addresses that are not valid (for example, 0. 0. 0. 0) and addresses for classes that are not supported (such as Class E). For entities that are not IP gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
- The type is int.
- The unit is datagrams.

#### IP Statistics Fragments Re-assembled/sec

- The rate that IP fragments are successfully re-assembled.
- The type is int.
- The unit is fragments/second.

#### IP Statistics Datagrams Received Header Errors

- The number of input datagrams discarded due to errors in their IP headers, such as bad checksums, version mismatch, other format errors, time-to-live exceeded, and errors discovered in processing their IP options.
- The type is int.
- The unit is datagrams.

#### IP Statistics Datagrams per Second

- The rate that IP datagrams are received from or sent to the interfaces, including those in error. Any forwarded datagrams are not included.
- The type is int.
- The unit is Datagrams/second.

#### IP Statistics Datagram Fragmentation Percentage

- Measure of datagram fragmentation. Calculated as  $100 * (\text{Fragments Received/sec} / \text{Datagrams Received/sec})$ .
- The type is int.
- The unit is percent.

#### IP Statistics Datagrams Outbound No Route

- The number of IP datagrams discarded because no route could be found to transmit them to their destination. This counter includes any packets counted in Datagrams Forwarded that meet this 'no route' criterion.
- The type is int.

- The unit is datagrams.

#### IP Statistics Fragment Re-assembly Failures

- The number of failures detected by the IP re-assembly algorithm (such as time out, errors, and so on). This is not necessarily a count of discarded IP fragments since some algorithms (notably RFC 815) can lose track of the number of fragments by combining them as they are received.
- The type is int.
- The unit is errors.

#### IP Statistics Fragmented Datagrams/sec

- The rate that datagrams are successfully fragmented at this entity.
- The type is int.
- The unit is datagrams/second.

#### IP Statistics Datagrams Forwarded/sec

- The rate of input datagrams for this entity was not their final IP destination. An attempt was made to find a route to forward the datagrams to their final destination. In entities that do not act as IP gateways, this rate includes only those packets that were Source-Routed via this entity, and the Source-Route option processing was successful.
- The type is int.
- The unit is datagrams/second.

#### IP Statistics Datagrams Received Unknown Protocol

- The number of locally-addressed datagrams received successfully, but discarded because of an unknown or unsupported protocol.
- The type is int.
- The unit is datagrams.

#### IP Statistics Datagrams Received/sec

- The rate that IP datagrams are received from the interfaces, including those in error.
- The type is int.
- The unit is datagrams/seconds.

#### IP Statistics Datagrams Outbound Discarded

- The number of output IP datagrams for which no problems were encountered to prevent their transmission to their destination, but which were discarded (for example, for lack of buffer space). This counter includes datagrams counted in Datagrams Forwarded, if any such packets met this (discretionary) discard criterion.
- The type is int.
- The unit is datagrams.

#### IP Statistics Fragments Created/sec

- The rate that datagrams are successfully fragmented at this entity.
- The type is int.
- The unit is datagrams/second.

#### IP Statistics Datagrams Received Discarded

- The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (for example, for lack of buffer space). This counter does not include any datagrams discarded while awaiting re-assembly.
- The type is int.
- The unit is datagrams.

#### IP Statistics Datagrams Sent/sec

- The rate that IP datagrams are supplied to IP for transmission by the local IP user-protocols (including ICMP). This counter does not include any datagrams counted in Datagrams Forwarded.
- The type is int.
- The unit is datagrams/second.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

### **Component: Active Tasks running on VCenter Server**

Use the Monitored Logs Report data set to monitor log settings that affect future log entries, such as the maximum log size and when old entries should be deleted. Monitored Logs Report is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### Monitored Logs Report Date Last Modified

- The date for when the log was last modified.
- The type is string.

##### Monitored Logs Report Event Log Type

- The unique integer that identifies the Event Log type. This value determines the which log table to call.
- The type is int.

##### Monitored Logs Report Time Last Modified

- The time for when the log was last modified.
- The type is string.

##### Monitored Logs Report Log Name

- Use this attribute to create a situation where you want to specify a specific log to monitor or exclude events written to a specific log.
- The type is string.

##### Monitored Logs Report Disk Drive Location

- A location on a disk drive. Valid format is a text string without case-sensitivity in the range from 1 to 256 characters. Do not use to create situations.
- The type is string.

##### Monitored Logs Report Date Time Last Modified

- The date and time when the file was last modified.
- The type is timestamp.

#### Monitored Logs Report Log Name (Unicode)

- Use this attribute to create a situation where you want to specify a specific log to monitor or exclude events written to a specific log.
- The type is string. This is a key dimension.

#### node

- The managed system name. The format should be *hostname : agent\_code* . Examples include spark:KNT or deux.raleigh.ibm.com:KNT.
- The type is string.

#### Monitored Logs Report Path (Unicode)

- A location on a disk drive. Valid format is a text string without case-sensitivity in the range from 1 to 392 bytes. Do not use to create situations.
- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### Retention Policy

- This attribute reflects the retention policy, expressed in days when the "Overwrite events older than" policy is in effect (for Windows 2003 and Windows XP only). Valid values include (0) Overwrite events as needed, (-1) Do not overwrite events, (-2) Archive the log when full. Do not overwrite events, (2147483647) Value\_Exceeds\_Maximum and (-2147483648) Value\_Exceeds\_Minimum.
- The type is int.

### Metrics

#### Record Count (Superseded)

- The number of records in a log.
- The type is int.
- The unit is records.

#### Current Size (Superseded)

- The current size of a specific log in bytes. This attribute reflects the size of the file on disk, and it might not match the value displayed by the event log viewer.
- The type is int.
- The unit is bytes.

#### Monitored Logs Report Maximum Size

- The maximum size of the log file in bytes. This attribute is the 64-bit version of Max Size.
- The type is double.
- The unit is bytes.

#### Monitored Logs Report Current Size

- The current size of a specific log in bytes. This attribute reflects the size of the file on disk, and it might not match the value displayed by the event log viewer. This attribute is the 64-bit version of Current Size.
- The type is double.
- The unit is bytes.

#### Maximum Size (Superseded)

- The maximum size of the log file in bytes.
- The type is int.
- The unit is bytes.

#### Row Number

- Row Number.
- The type is int.
- The unit is row.

#### Monitored Logs Report Record Count

- The number of records in a log. This attribute is the 64-bit version of Record Count.
- The type is double.
- The unit is records.

#### Monitored Logs Report Percent of Log Used

- The percentage of the log used.
- The type is int.
- The unit is percent.

### **Component: Active Tasks running on VCenter Server**

Network Interface Object.

#### **Dimensions**

##### DHCP

- Indicates whether or not the dynamic host configuration protocol (DHCP) is enabled for this adapter.
- The type is int.

##### Description

- The description of the Network Interface.
- The type is string.

##### Network Interface Instance

- The instance name of the Network Interface object (Connection Name). The valid format is a text string of up to 128 characters.
- The type is string. This is a key dimension.

##### IPv4 Address

- The IPv4 interface address.
- The type is string.

#### Network Interface Instance (Unicode)

- The instance name of the Network Interface object (Connection Name) in unicode.
- The type is string.

#### IPv6 Link Local Address

- The IPv6 Link Local interface address.
- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### IPv6 Global Address

- The IPv6 Global interface address. Note that the value No\_DNS\_Entry is a valid value.
- The type is string.

#### Friendly Name

- The friendly name of the Network Interface.
- The type is string. This is a key dimension.

#### node

- The managed system name.
- The type is string.

### **Metrics**

#### Network Interface Bytes Sent/sec

- The rate that bytes are sent on the interface, including framing characters.
- The type is double.
- The unit is bytes/second.

#### Packets Received Discarded

- The number of inbound packets that were chosen to be discarded even though no errors had been detected to prevent their being delivered to a higher-layer protocol.
- The type is double.
- The unit is packets.

#### Network Interface Bytes Received/sec

- The rate that bytes are received on the interface, including framing characters.
- The type is double.
- The unit is bytes/second.

#### Output Queue Length kPackets

- The length of the output packet queue in packets (in thousands). If this is longer than 2, delays are being experienced and the bottleneck must be found and eliminated, if possible.
- The type is double.
- The unit is kilopackets.



#### Packets Outbound Errors

- The number of outbound packets that could not be transmitted because of errors.
- The type is double.
- The unit is errors.

#### Packets Received Errors

- The number of inbound packets that contained errors preventing them from being delivered to a higher-layer protocol.
- The type is double.
- The unit is packets.

#### Packets Received/sec

- The rate that packets are received on the network interface.
- The type is double.
- The unit is packets/second.

#### Packets Sent Unicast/sec

- The rate that packets are requested to be transmitted to subnet-unicast addresses by higher-layer protocols. The rate includes the packets that were discarded or not sent.
- The type is double.
- The unit is packets/second.

#### Output Queue Length

- The length of the output packet queue (in packets). If this is longer than 2, delays are being experienced and the bottleneck must be found and eliminated, if possible.
- The type is double.
- The unit is packets.

#### Current Bandwidth

- An estimate of the current bandwidth for the interface in bits per second (bps). For interfaces that do not vary in bandwidth, or for those where no accurate estimation can be made, this value is the nominal bandwidth.
- The type is double.
- The unit is bits/second.

#### Packets Sent/sec

- The rate that packets are sent on the network interface.
- The type is double.
- The unit is packets/second.

#### Packets Sent Non-Unicast/sec

- The rate that packets are requested to be transmitted to non-unicast, that is, subnet broadcast or subnet multicast, addresses by higher-layer protocols. The rate includes the packets that were discarded or not sent.
- The type is double.
- The unit is packets/second.

#### Packets Received Non-Unicast/sec

- The rate that non-unicast, that is, subnet broadcast or subnet multicast packets, are delivered to a higher-layer protocol.
- The type is double.
- The unit is packets/second.

#### Network Interface Bytes Total/sec

- The rate that bytes are sent and received on the interface,.
- The type is double.
- The unit is bytes/second.

#### Bandwidth Utilization Percentage

- Measure of network interface bandwidth utilization. Calculated as  $100 * ((8 * \text{Bytes Total/sec}) / \text{Current Bandwidth})$ .
- The type is int.
- The unit is percent.

#### Packets/sec

- The rate that packets are sent and received per second on the network interface.
- The type is double.
- The unit is packets/second.

#### Packets Received Unicast/sec

- The rate that (subnet) unicast packets are delivered to a higher-layer protocol.
- The type is double.
- The unit is packets/second.

#### Packets Outbound Discarded

- The number of outbound packets that were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space.
- The type is double.
- The unit is packets.

#### Packets Received Unknown

- The number of packets received via the interface that were discarded because of an unknown or unsupported protocol.
- The type is double.
- The unit is packets.

### **Component: Active Tasks running on VCenter Server**

Use the MSMQ Queue data set to monitor MSMQ (Microsoft Message Queue) statistics. MSMQ Queue is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### MSMQ Queue Instance

- The instance name of the queue. For example, LMAIL.
- The type is string. This is a key dimension.

#### MSMQ Queue Instance (Unicode)

- The instance name of the queue. Queue\_is\_Inactive is displayed for this attribute if there are no messages in the queue.
- The type is string. This is a key dimension.

#### node

- The managed system name. The form should be *hostname : agent\_code* . Examples include spark:KNT or deux.raleigh.ibm.com:KNT.
- The type is string.

### Metrics

#### MSMQ Bytes in Journal Queue

- The total number of bytes that currently reside in the journal queue. For the Computer Queues instance, this represents the computer journal queue.
- The type is int.
- The unit is bytes.

#### MSMQ Messages in Queue

- The total number of messages that currently reside in the queue. For the Computer Queues instance, this represents the dead letter queue.
- The type is int.
- The unit is messages.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### MSMQ Bytes in Queue

- The total number of bytes that currently reside in the queue. For the Computer Queues instance, this represents the dead letter queue.
- The type is int.
- The unit is bytes.

#### MSMQ Messages in Journal Queue

- The total number of messages that currently reside in the journal queue. For the Computer Queues instance, this represents the dead letter queue.
- The type is int.
- The unit is messages.

## Component: Active Tasks running on VCenter Server

FTP Service data set to create situations that monitor traffic and connection activity for an FTP (File Transfer Protocol) server. FTP Service is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

### Dimensions

node

- The managed system name.
- The type is string.

FTP Service File Transfer Protocol Site

- Name of File Transfer Protocol site.
- The type is string. This is a key dimension.

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

### Metrics

FTP Service Maximum Anonymous Users

- The maximum number of users who established concurrent anonymous connections using the FTP service (since service startup).
- The type is int.
- The unit is users.

FTP Service Total Files Received

- The total number of files received by the FTP service.
- The type is int.
- The unit is files.

FTP Service Maximum NonAnonymous Users

- The maximum number of users who established concurrent non-anonymous connections using the FTP service (since service startup).
- The type is int.
- The unit is users.

FTP Service Bytes Sent/sec

- The rate that data bytes are sent by the FTP service.
- The type is int.
- The unit is bytes/second.

FTP Service Bytes Total/sec

- The sum of Bytes Sent/sec and Bytes Received/sec. This is the total rate of bytes transferred by the FTP service.
- The type is int.
- The unit is bytes/second.

FTP Service Current NonAnonymous Users

- The number of users who currently have a non-anonymous connection using the FTP service.
- The type is int.
- The unit is users.

#### FTP Service Current Connections

- The current number of connections established with the FTP service.
- The type is int.
- The unit is connections.

#### FTP Service Total Logon Attempts

- The number of logons that have been attempted using the FTP service (since service startup).
- The type is int.
- The unit is attempts.

#### FTP Service Maximum Connections

- The maximum number of simultaneous connections established with the FTP service.
- The type is int.
- The unit is connections.

#### FTP Service Bytes Received/sec

- The rate that data bytes are received by the FTP service.
- The type is int.
- The unit is bytes/second.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### FTP Service Current Anonymous Users

- The number of users who currently have an anonymous connection using the FTP service.
- The type is int.
- The unit is users.

#### FTP Service Total Files Sent

- The total number of files sent by the FTP service since service startup.
- The type is int.
- The unit is files.

#### FTP Service Total Files Transferred

- The sum of Files Sent and Files Received. This is the total number of files transferred by the FTP service since service startup.
- The type is int.
- The unit is files.

#### FTP Service Total Connection Attempts

- The number of connections that have been attempted using the FTP service (since service startup). This counter is for all instances listed.
- The type is int.
- The unit is connections.

#### FTP Service Total NonAnonymous Users

- The total number of users who established a non-anonymous connection with the FTP service (since service startup).
- The type is int.
- The unit is users.

#### FTP Service Total Anonymous Users

- The total number of users who established an anonymous connection with the FTP service (since service startup).
- The type is int.
- The unit is users.

### **Component: Mount Points**

Information about the mounted volumes.

#### **Dimensions**

##### Mount Point

- The mount point where the volume is mounted.
- The type is string.

##### Mount Point State

- The mount point state.
- The type is string.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### Mount Point Drive Name

- The drive name.
- The type is string.

##### Mounted Volume Name

- The volume name.
- The type is string.

##### node

- The managed system name.
- The type is string.

### **Component: DNS Zone Transfer**

Information about DNS (Domain Name Server) server transaction activity and performance.

## Dimensions

### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

### DNS Zone Transfer node

- The managed system name.
- The type is string.

### DNS Zone Transfer DNS Parameter

- A special one for the front end to use as a column header.
- The type is string.

### DNS Zone Transfer DNS Value

- A special one for the front end to use as a column header.
- The type is string.

## Metrics

### DNS Zone Transfer Request Received

- The total number of zone transfer requests received by the master DNS server.
- The type is int.
- The unit is requests.

### DNS Zone Transfer DNS AXFR Success Sent

- The total number of successful full zone transfers of the master DNS server.
- The type is int.
- The unit is transfers.

### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

### DNS Zone Transfer DNS AXFR Request Received

- The total number of full zone transfer requests received by the secondary DNS server.
- The type is int.
- The unit is requests.

### DNS Zone Transfer DNS IXFR Request Received

- The total number of incremental zone transfer requests received by the master DNS server.
- The type is int.
- The unit is requests.

### DNS Zone Transfer DNS Notify Sent

- The total number of notifies sent by the master DNS server.

- The type is int.
- The unit is notifies.

#### DNS Zone Transfer DNS IXFR Success Received

- The total number of successful incremental zone transfers received by the secondary DNS server.
- The type is int.
- The unit is transfers.

#### DNS Zone Transfer Failure

- The total number of failed zone transfers of the master DNS server.
- The type is int.
- The unit is transfers.

#### DNS Zone Transfer DNS IXFR Response Received

- The total number of incremental zone transfer responses received by the secondary DNS server.
- The type is int.
- The unit is responses.

#### DNS Zone Transfer DNS IXFR Success Sent

- The total number of successful incremental zone transfers of the master DNS server.
- The type is int.
- The unit is transfers.

#### DNS Zone Transfer DNS AXFR Success Received

- The total number of successful full zone transfers received by the secondary DNS server.
- The type is int.
- The unit is transfers.

#### DNS Zone Transfer SOA Request Sent

- The total number of zone transfer SOA requests sent by the secondary DNS server.
- The type is int.
- The unit is requests.

#### DNS Zone Transfer DNS IXFR Request Sent

- The total number of incremental zone transfer requests sent by the secondary DNS server.
- The type is int.
- The unit is requests.

#### DNS Zone Transfer DNS AXFR Request Sent

- The total number of full zone transfer requests sent by the secondary DNS server.
- The type is int.
- The unit is requests.

#### DNS Zone Transfer DNS AXFR Response Received

- The total number of full zone transfer responses received by the secondary DNS server. Valid values are positive integers in the range 0 to 2147483647 and can include the use of the \*AVG,



\*MIN, \*MAX, or \*SUM functions. Note: the value 2147483647 indicates Value\_Exceeds\_Maximum and the value -2147483648 indicates Value\_Exceeds\_Minimum.

- The type is int.
- The unit is responses.

#### DNS Zone Transfer DNS IXFR TCP Success Received

- The total number of successful TCP incremental zone transfers received by the secondary DNS server.
- The type is int.
- The unit is transfers.

#### DNS Zone Transfer DNS Notify Received

- The total number of notifies received by the secondary DNS server.
- The type is int.
- The unit is notifies.

#### DNS Zone Transfer DNS IXFR UDP Success Received

- The total number of successful UDP incremental zone transfers received by the secondary DNS server.
- The type is int.
- The unit is transfers.

#### DNS Zone Transfer Success

- The total number of successful zone transfers of the master DNS server.
- The type is int.
- The unit is transfers.

### **Component: Active Tasks running on VCenter Server**

Use the Devices data set to create situations to obtain status and configuration information about all of the devices installed on the Windows Server. Devices is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### Devices Internal Device Name

- The internal name of the device in the Service Control Manager database. The maximum size of the text string is 256 bytes, but here it is truncated to 64 bytes.
- The type is string. This is a key dimension.

##### Device Dependency

- The name of a device or load order group that must start before the given device can start. If there are no dependencies for the given device, this field is blank. For example, +SCSI miniport indicates the name of a device that must start before the given device can start.
- The type is string. This is a key dimension.

##### Devices Display Name (Unicode)

- The name of the driver as it is displayed in the Service Control Manager applet in UTF8.
- The type is string.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### Devices Path to Device Executable

- The fully qualified path to the device binary executable. For example, SystemRootSystem32driversafd.sys indicates the path to the device binary executable.
- The type is string.

#### Devices Current State

- The current state of the device, which can be Stopped, Start Pending, Stop Pending, Running, Continue Pending, Paused Pending, or Paused. For example, Running indicates that the device is currently running.
- The type is string.

#### Devices Start Type

- Specifies how to start the device, including Automatic, Manual, Disabled, Boot, System, and Unknown.
- The type is string.

#### Devices Driver Object Name

- Specifies an object name. If the service is of type WIN32, this is the account name that the service uses to log on when it runs. If the service is type Kernel Driver or File System Driver, this is the Windows Server's driver object name that the I/O Manager uses to load the device driver.
- The type is string.

#### Devices Binary Path (Unicode)

- The fully qualified path to the device binary executable in UTF8.
- The type is string.

#### Devices Load Order Group

- The name of the load ordering group of which this device is a member. Devices can be placed in groups so other devices can have dependencies on a group of devices. If the device is not in a load ordering group, this field is blank. For example, SCSI CDROM Class is an example of a load order group.
- The type is string.

#### node

- The managed system name. The format should be *hostname : agent\_code* . Examples include spark:KNT or deux.raleigh.ibm.com:KNT.
- The type is string.

#### Devices Name of the Driver

- The name of the driver as it is displayed in the Windows Service Control Manager applet. Valid format is a text string of up to 64 characters. For example, Cdrom is an example of a driver name.
- The type is string.

### Metrics

#### Row Number

- Row Number.

- The type is int.
- The unit is row.

### **Component: Active Tasks running on VCenter Server**

Use the Network Segment data set to monitor utilization and traffic statistics for data in a network segment. Data for this report is from the Microsoft Network Monitor. Network Segment is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set. Network Segment attributes are only supported for Windows NT and earlier. They exist in this release for backward compatibility only.

#### **Dimensions**

##### Collection Time

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### Network Segment Instance

- The instance name of the Network Interface object (Connection Name).
- The type is string. This is a key dimension.

##### node

- The managed system name.
- The type is string.

#### **Metrics**

##### Network Segment % Network Utilization

- The percentage of network bandwidth in use of this network segment.
- The type is int.
- The unit is percent.

##### Network Segment Multicast Frames Received/sec

- The number of multicast frames received per second on this network segment. Valid format is a text string of up to 64 characters.
- The type is int.
- The unit is frames/second.

##### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

##### Network Segment % Broadcast Frames

- The percentage of network bandwidth that is made up of broadcast traffic on this network segment. Valid values are positive integers in the range 0 to 100 (expressing a percentage).
- The type is int.
- The unit is percent.

##### Network Segment Total Frames Received/sec

- The total number of frames received per second on this network segment.
- The type is int.
- The unit is frames/second.

#### Network Segment Total Bytes Received/sec

- The number of bytes received per second on this network segment.
- The type is int.
- The unit is bytes/second.

#### Network Segment % Multicast Frames

- The percentage of network bandwidth that is made up of multicast traffic on this network segment. Valid values are positive integers in the range 0 to 100 (expressing a percentage).
- The type is int.
- The unit is percent.

#### Network Segment Broadcast Frames Received/sec

- The number of broadcast frames received per second on this network segment.
- The type is int.
- The unit is frames/second.

### **Component: HTTP Service**

Information about traffic and connection activity for an HTTP (HyperText Transport Protocol) server, such as the current connections, the bytes received per second, and the total anonymous users connected.

#### **Dimensions**

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### HTTP Service Value

- A special one for the front end to use as a column header.
- The type is string.

##### HTTP Service Parameter

- A special one for the front end to use as a column header.
- The type is string.

##### node

- The managed system name.
- The type is string.

#### **Metrics**

##### HTTP Service Logon Attempts

- The number of logon attempts that have been made by the HTTP server.
- The type is int.
- The unit is attempts.

#### HTTP Service Sent and Received Files

- The sum of files sent and files received. This is the total number of files transferred by the HTTP server.
- The type is int.
- The unit is files.

#### HTTP Service Current ISAPI Extension Requests

- The number of ISAPI (Internet Server API) Extension requests that are simultaneously being processed by the HTTP server.
- The type is int.
- The unit is requests.

#### HTTP Service Maximum CGI Requests

- The maximum number of CGI requests that have been simultaneously processed by the HTTP server. This includes WAIS index queries.
- The type is int.
- The unit is requests.

#### HTTP Service HTTP Requests

- The number of HTTP requests being handled per second.
- The type is int.
- The unit is connections/second.

#### HTTP Service Bytes Sent/sec

- The rate that data bytes are sent by the HTTP server.
- The type is int.
- The unit is bytes/second.

#### HTTP Service CGI Requests

- Common Gateway Interface (CGI) requests are custom gateway executables (exe). An administrator can install these executables to add forms processing or other dynamic data sources.
- The type is int.
- The unit is requests.

#### HTTP Service ISAPI Extension Requests

- The number of ISAPI Extension requests. ISAPI Extension requests are custom gateway Dynamic Link Libraries (.dll) an administrator can install to add forms processing to other dynamic data sources.
- The type is int.
- The unit is requests.

#### HTTP Service Connection Attempts

- The number of connection attempts that have been made to the HTTP server.
- The type is int.
- The unit is attempts.

#### HTTP Service Current Anonymous Users

- The number of anonymous users currently connected to the HTTP server.
- The type is int.
- The unit is users.

#### HTTP Service Get Requests

- The number of HTTP requests using the GET method. Get requests are generally used for basic file retrievals or image maps, though they can be used with forms.
- The type is int.
- The unit is requests.

#### HTTP Service Post Requests

- The number of HTTP requests using the POST method. Post requests are generally used for forms or gateway requests.
- The type is int.
- The unit is requests.

#### HTTP Service Current NonAnonymous Users

- The number of non-anonymous users currently connected to the HTTP server.
- The type is int.
- The unit is users.

#### HTTP Service Maximum ISAPI Extension Requests

- The maximum number of ISAPI Extension requests that have been simultaneously processed by the HTTP server.
- The type is int.
- The unit is requests.

#### HTTP Service Total Anonymous Users

- The total number of anonymous users that have ever connected to the HTTP server.
- The type is int.
- The unit is users.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anynother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### HTTP Service Other Requests

- The number of HTTP requests that are not GET, POST, or HEAD methods. These might include PUT, DELETE, LINK, or other methods supported by gateway applications.
- The type is int.
- The unit is requests.

#### HTTP Service Current Connections

- The current number of connections to the HTTP server.
- The type is int.

- The unit is connections.

#### HTTP Service FTP Server Files Sent

- The total number of files sent by the HTTP server.
- The type is int.
- The unit is files.

#### HTTP Service Bytes Received/sec

- The rate that data bytes are received by the HTTP server.
- The type is int.
- The unit is bytes/second.

#### HTTP Service Total NonAnonymous Users

- The total number of non-anonymous users that have ever connected to the HTTP server.
- The type is int.
- The unit is users.

#### HTTP Service Maximum Connections

- The maximum number of simultaneous connections to the HTTP Server.
- The type is int.
- The unit is connections.

#### HTTP Service Current CGI Requests

- The current number of CGI requests that are simultaneously being processed by the HTTP server. This includes WAIS index queries.
- The type is int.
- The unit is requests.

#### HTTP Service Head Requests

- The number of HTTP requests using the HEAD method. Head requests generally indicate a client is querying the state of a document they already have to see if it needs to be refreshed.
- The type is int.
- The unit is requests.

#### HTTP Service HTTP Server Files Received

- The total number of files received by the HTTP server.
- The type is int.
- The unit is files.

#### HTTP Service Not Found Errors

- The number of requests that could not be satisfied by the server because the requested document could not be found. These are generally reported as an HTTP 404 error code to the client.
- The type is int.
- The unit is requests.

#### HTTP Service Bytes Total/sec

- The sum of bytes sent per second and bytes received per second. This is the total rate of bytes transferred by the HTTP server.
- The type is int.
- The unit is bytes/second.

#### HTTP Service Maximum Anonymous Users

- The maximum number of anonymous users simultaneously connected to the HTTP server.
- The type is int.
- The unit is users.

#### HTTP Service Maximum NonAnonymous Users

- The maximum number of non-anonymous users simultaneously connected to the HTTP server.
- The type is int.
- The unit is users.

### **Component: Active Tasks running on VCenter Server**

Use the RAS Port data set to monitor Remote Access Service Port activity. RAS Port is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### Remote Access Service Port Instance

- The instance name of the queue. Valid format is a text string of up to 64 characters. For example, tAM1.
- The type is string. This is a key dimension.

##### node

- The managed system name.
- The type is string.

#### **Metrics**

##### Remote Access Svc Port Bytes Received/sec

- The number of bytes received per second for this connection.
- The type is int.
- The unit is bytes/second.

##### Remote Access Service Port Bytes Transmitted/sec

- The number of bytes transmitted per second for this connection.
- The type is int.
- The unit is bytes/second.

##### Remote Access Service Port Total Errors/sec

- The total number of CRC, Timeout, Serial Overrun, Alignment, and Buffer Overrun errors per second.
- The type is int.



- The unit is errors/second.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### Remote Access Service Port Serial Overrun Errors

- The total number of serial overrun errors for this connection. Serial overrun errors occur when the hardware cannot handle the rate at which data is received.
- The type is int.
- The unit is errors.

#### Remote Access Service Port Total Errors

- The total number of CRC, Timeout, Serial Overrun, Alignment, and Buffer Overrun errors for this connection.
- The type is int.
- The unit is errors.

#### Remote Access Service Port Frames Received/sec

- The number of data frames received per second for this connection.
- The type is int.
- The unit is frames/second.

#### Remote Access Svc Port Alignment Errors

- The total number of alignment errors for this connection. Alignment errors occur when a byte received is different from the byte expected.
- The type is int.
- The unit is errors.

#### Remote Access Svc Port Buffer Overrun Errors

- The total number of buffer overrun errors for this connection. Buffer overrun errors occur when the software cannot handle the rate at which data is received.
- The type is int.
- The unit is errors.

#### Remote Access Service Port Bytes Transmitted

- The total number of bytes transmitted for this connection.
- The type is int.
- The unit is bytes.

#### Remote Access Service Port Frames Transmitted

- The total number of data frames transmitted for this connection.
- The type is int.
- The unit is frames.

#### Remote Access Service Port CRC Errors

- The total number of CRC errors for this connection. CRC errors occur when the frame received contains erroneous data.
- The type is int.
- The unit is errors.

#### Remote Access Service Port Timeout Errors

- The total number of Timeout Errors for this connection. Timeout errors occur when an expected is not received in time.
- The type is int.
- The unit is errors.

#### Remote Access Service Port Frames Transmitted/sec

- The number of data frames transmitted per second for this connection.
- The type is int.
- The unit is frames/second.

#### Remote Access Service Port Percent Compression In

- The compression ratio for bytes being received.
- The type is int.
- The unit is percent.

#### Remote Access Service Port Percent Compression Out

- The compression ratio for bytes being transmitted.
- The type is int.
- The unit is percent.

#### Remote Access Svc Port Bytes Received

- The total number of bytes received for this connection.
- The type is int.
- The unit is bytes.

#### Remote Access Service Port Frames Received

- The total number of data frames received for this connection.
- The type is int.
- The unit is frames.

### **Component: Server**

Information about connections and throughput between the local computer (Server/Redirector) and the network.

#### **Dimensions**

##### Server Parameter

- A special one for the front end to use as a column header.
- The type is string.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.

- The type is timestamp.

node

- The managed system name.
- The type is string.

Server Value

- A special one for the front end to use as a column header.
- The type is string.

## Metrics

Server Sessions

- The number of sessions currently active in the server. Indicates current server activity.
- The type is int.
- The unit is sessions.

Server Total Ended Sessions (Superseded)

- Total number of sessions that have ended. Calculated as the sum of Sessions Errored Out, Sessions Forced Off, Sessions Logged Off, and Sessions Timed Out.
- The type is int.
- The unit is sessions.

Server Files Opened Total

- The number of successful open attempts performed by the server on behalf of clients. Useful in determining the amount of file I/O, determining overhead for path-based operations, and for determining the effectiveness of open locks.
- The type is int.
- The unit is opens.

Server Logon Total

- Includes all interactive logons, network logons, service logons, successful logons, and failed logons since the system is last rebooted.
- The type is int.
- The unit is logons.

Server Errors Granted Access

- The number of times accesses to files opened successfully were denied. Can indicate attempts to access files without proper access authorization.
- The type is int.
- The unit is errors.

Server Work Item Shortages

- The number of times STATUS\_DATA\_NOT\_ACCEPTED was returned at receive indication time. This occurs when no work item is available or can be allocated to service the incoming request. Indicates whether the InitWorkItems or MaxWorkItems parameters might need to be adjusted.
- The type is int.
- The unit is count.

#### Server Bytes Total/sec

- The number of bytes the server has sent to and received from the network. This value provides an overall indication of how busy the server is. This attribute is the 64-bit version of Bytes Total/sec.
- The type is double.
- The unit is bytes/second.

#### Server Pool Paged Bytes

- The number of bytes of pageable computer memory the server is currently using. Can help in determining good values for the MaxPagedMemoryUsage parameter.
- The type is int.
- The unit is bytes.

#### Server Bytes Received/sec

- The number of bytes the server has received from the network. Indicates how busy the server is. This attribute is the 64-bit version of Bytes Received/sec.
- The type is double.
- The unit is bytes/second.

#### Server Bytes Total/sec (Superseded)

- The number of bytes the server has sent to and received from the network. This value provides an overall indication of how busy the server is.
- The type is int.
- The unit is bytes/second.

#### Server Pool Paged Failures

- The number of times allocations from paged pool have failed. Indicates that the computer's physical memory or paging file are too small.
- The type is int.
- The unit is errors.

#### Server Bytes Received/sec (Superseded)

- The number of bytes the server has received from the network. Indicates how busy the server is.
- The type is int.
- The unit is bytes/second.

#### Server File Directory Searches

- The number of searches for files currently active in the server. Indicates current server activity.
- The type is int.
- The unit is searches.

#### Server Context Blocks Queued/sec

- The rate at which work context blocks had to be placed on the FSP queue of the server to await server action.
- The type is int.
- The unit is blocks/second.

#### Server Errors System

- The number of times an internal Server Error was detected. Unexpected errors usually indicate a problem with the Server.
- The type is int.
- The unit is errors.

#### Server Bytes Transmitted/sec (Superseded)

- The number of bytes the server has sent on the network. Indicates how busy the server is.
- The type is int.
- The unit is bytes/second.

#### Row Number

- Row Number.
- The type is int.
- The unit is row.

#### Server Error Session Percent

- Percentage of total sessions that ended due to errors.
- The type is int.
- The unit is percent.

#### Server High % Bytes/Sec

- The percentage of network card bandwidth used by the server service.
- The type is int.
- The unit is percent/second.

#### Server Errors Logon

- The number of failed logon attempts to the server. Can indicate whether password guessing programs are being used to crack the security on the server.
- The type is int.
- The unit is errors.

#### Server Sessions Timed Out

- The number of sessions that have been closed due to their idle time exceeding the AutoDisconnect parameter for the server. Shows whether the AutoDisconnect setting is helping to conserve resources.
- The type is int.
- The unit is sessions.

#### Server Pool Nonpaged Bytes

- The number of bytes of non-pageable computer memory the server is using. This value is useful for determining the values of the MaxNonpagedMemoryUsage value entry in the Windows NT Registry.
- The type is int.
- The unit is bytes.

#### Server Pool Nonpaged Peak

- The maximum number of bytes of nonpaged pool the server has had in use at any one point. Indicates how much physical memory the computer should have.

- The type is int.
- The unit is bytes.

#### Server Bytes Transmitted/sec

- The number of bytes the server has sent on the network. Indicates how busy the server is. This attribute is the 64-bit version of Bytes\_Transmitted/sec.
- The type is double.
- The unit is bytes/second.

#### Server Files Open

- The number of files currently opened in the server. Indicates current server activity.
- The type is int.
- The unit is opens.

#### Server Pool Paged Peak

- The maximum number of bytes of paged pool the server has had allocated. Indicates the proper sizes of the Page File(s) and physical memory.
- The type is int.
- The unit is bytes.

#### Server Total Ended Sessions

- Total number of sessions that have ended. Calculated as the sum of Sessions Errored Out, Sessions Forced Off, Sessions Logged Off, and Sessions Timed Out. This attribute is the 64-bit version of Total Ended Sessions.
- The type is double.
- The unit is sessions.

#### Server Sessions Forced Off

- The number of sessions that have been forced to logoff. Can indicate how many sessions were forced to logoff due to logon time constraints.
- The type is int.
- The unit is sessions.

#### Server Sessions Logged Off

- The number of sessions that have terminated normally. Useful in interpreting the Sessions Times Out and Sessions Errored Out statistics, allows percentage calculations.
- The type is int.
- The unit is sessions.

#### Server Errors Access Permissions

- The number of times opens on behalf of clients have failed with STATUS\_ACCESS\_DENIED. Can indicate whether somebody is randomly attempting to access files in hopes of getting at something that was not properly protected.
- The type is int.
- The unit is errors.

#### Server Logons sec

- The rate of all server logons.

- The type is int.
- The unit is logons/second.

#### Server Blocking Requests Rejected

- The number of times the server has rejected blocking SMBs due to insufficient count of free work items. Indicates whether the MaxWorkItem or MinFreeWorkItems server parameters might need to be adjusted.
- The type is int.
- The unit is rejections.

#### Server Sessions Errored Out

- The number of sessions that have been closed due to unexpected error conditions or sessions that have reached the autodisconnect timeout and have been disconnected normally.
- The type is int.
- The unit is sessions.

#### Server Pool Nonpaged Failures

- The number of times allocations from nonpaged pool have failed. Indicates that the computer's physical memory is too small.
- The type is int.
- The unit is errors.

### **Component: Windows Memory**

Information about real and virtual memory. Real memory is allocated in units of pages. Virtual memory may exceed real memory size, causing page traffic as virtual pages are moved between disk and real memory.

#### **Dimensions**

node

- The managed system name.
- The type is string.

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

Memory Total Memory Size (MB)

- The number of megabytes of installed random access memory (RAM) in the computer.
- The type is int.

Memory Total Memory Size (Bytes)

- The number of bytes of installed random access memory (RAM) in the computer.
- The type is double.

Memory Total Memory Size (KB)

- The number of kilobytes of installed random access memory (RAM) in the computer. Note: -1 indicates Unknown and 9223372036854775807 indicates Value\_Exceeds\_Maximum.
- The type is double.

## Metrics

### Memory Pool Paged Bytes

- The number of bytes in the paged pool area, a system memory area where space is acquired by operating system components as they accomplish their appointed tasks. Paged pool pages can be paged out to the paging file when not accessed by the system for sustained periods of time.
- The type is double.
- The unit is bytes.

### Memory Cache Bytes Peak

- The maximum number of bytes used by the system cache. The system cache is used to buffer data retrieved from disk or LAN. The system cache uses memory not in use by active processes in the computer.
- The type is double.
- The unit is bytes.

### Memory Available Usage Percentage

- The percent of bytes of real memory available. Calculated as  $100 * (\text{Available Bytes} / \text{Total Memory Bytes})$ .
- The type is int.
- The unit is percent.

### Memory Available kBytes

- The amount of available real memory, in KBs.
- The type is double.
- The unit is kilobytes.

### Memory Commit Avail kBytes

- The number of KBs until the commit limit is reached. Calculated as  $\text{Commit Limit kB} - \text{Committed kBytes}$ .
- The type is double.
- The unit is kilobytes.

### Memory Cache Bytes

- The number of bytes currently in use by the system cache. The system cache is used to buffer data retrieved from disk or LAN. The system cache uses memory not in use by active processes in the computer.
- The type is double.
- The unit is bytes.

### Memory Commit Limit (Bytes)

- The number of bytes of virtual memory that can be committed on a system without extending the paging files. Commit limit is the size (in bytes) of virtual memory that can be committed without having to extend the paging files. If the paging files can be extended, this is a soft limit.
- The type is double.
- The unit is bytes.

### Memory System Code Total Bytes

- The number of bytes of pageable operating system code that is currently in virtual memory.



- The type is double.
- The unit is bytes.

#### Memory Page Faults/sec

- The number of page faults in the processor. Note that page faults occur when a process refers to a page that is not in its Working Set in main memory. A Page Fault will not cause the page to be fetched from disk if that page is on the standby list, and hence already in main memory, or if it is in use by another process with whom the page is shared.
- The type is int.
- The unit is faults/second.

#### Memory Pool Paged Resident Bytes

- The current size of the paged pool. The paged pool is an area of physical memory acquired by the operating system for objects that can be paged out to the paging file when not accessed by the system for sustained periods of time.
- The type is double.
- The unit is bytes.

#### Memory System Driver Resident Bytes

- The number of bytes of pageable physical memory being used by device drivers.
- The type is double.
- The unit is bytes.

#### Memory Pages/sec

- The number of pages read from the disk or written to the disk to resolve memory references to pages that were not in memory at the time of the reference. This is the sum of Pages Input/Sec and Pages Output/Sec attributes. This counter includes paging traffic on behalf of the system cache to access file data for applications. This is the primary counter to observe if you are concerned about excessive memory pressure (that is, thrashing), and the excessive paging that may result.
- The type is int.
- The unit is pages/second.

#### Memory Page Reads/sec

- The number of times the disk was read to retrieve pages of virtual memory necessary to resolve page faults. Multiple pages can be read during a disk read operation.
- The type is int.
- The unit is reads/second.

#### Memory Cache Usage Percentage

- The percent of bytes of real memory allocated to the system cache. Calculated as  $100 * (\text{Cache Bytes} / \text{Total Memory Bytes})$ .
- The type is int.
- The unit is percent.

#### Memory Usage Percentage

- The percent of bytes of real memory in use. Calculated as  $100 * (\text{Committed Bytes} / \text{Total Memory Bytes})$ .
- The type is int.

- The unit is percent.

#### Memory % Committed Bytes In Use

- The ratio of Committed Bytes to the Commit Limit. Committed memory is the physical memory in use for which space has been reserved in the paging file should it need to be written to disk. The commit limit is determined by the size of the paging file. If the paging file is enlarged, commit limit is enlarged, the commit limit increases, and the ratio is reduced. This counter displays the current percentage value only; it is not an average.
- The type is int.
- The unit is percent.

#### Memory Page Writes/sec

- The count of the number of times pages have been written to the disk because they were changed since last retrieved. Each such write operation may transfer a number of pages.
- The type is int.
- The unit is writes/second.

#### Memory Pool Nonpaged Allocs

- The number of calls to allocate space in the system nonpaged pool. Nonpaged pool is a system memory area where space is acquired by operating system components as they accomplish their appointed tasks. nonpaged pool pages cannot be paged out to the paging file, but instead remain in main memory as long as they are allocated.
- The type is int.
- The unit is calls.

#### Memory Working Set Total kBytes

- The number of KBs of real memory allocated to currently running processes.
- The type is double.
- The unit is kilobytes.

#### Memory Pages Output/sec

- The count of the number of pages that are written to disk because the pages have been modified in main memory.
- The type is int.
- The unit is pages/second.

#### Memory Committed kBytes

- The number of KBs of virtual memory that have been committed on a system.
- The type is double.
- The unit is kilobytes.

#### Memory System Driver Total Bytes

- The number of bytes of pageable virtual memory that is currently being used by device drivers.
- The type is double.
- The unit is bytes.

#### Memory Available Bytes

- The size of the virtual memory currently on the Zeroed, Free and Standby lists. Zeroed and Free memory is ready for use, with Zeroed memory cleared to zeros. Standby memory is memory

removed from a process Working Set but still available. Notice that this is an instantaneous count, not an average over the time interval.

- The type is double.
- The unit is bytes.

#### Memory Cache kBytes

- The amount of cache memory, in KBs, the system is currently using.
- The type is double.
- The unit is kilobytes.

#### Memory Pool Nonpaged kBytes

- The number of KBs in the nonpaged pool area of memory.
- The type is double.
- The unit is kilobytes.

#### Memory Pool Paged kBytes

- The number of KBs in the paged pool area of memory.
- The type is double.
- The unit is kilobytes.

#### Memory Demand Zero Faults/sec

- The rate at which a zeroed page is required to satisfy the fault. Zeroed pages are pages emptied of previously stored data and filled with zeros. These are a security feature of Windows that prevent processes from seeing data stored by earlier processes that used the memory space. Windows maintains a list of zeroed pages to accelerate this process. This counter shows the number of faults, without regard to the number of pages retrieved to satisfy the fault. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
- The type is int.
- The unit is faults/second.

#### Memory Committed Bytes

- The number of bytes of virtual memory that have been committed on a system. Committed memory must have disk storage available, or must be assured never to need disk storage (because main memory is large enough to hold it). Notice that this is an instantaneous count, not an average over the time interval.
- The type is double.
- The unit is bytes.

#### Memory Transition Faults/sec

- The number of page faults resolved by recovering pages that were in transition, that is, being written to disk at the time of the page fault. The pages were recovered without additional disk activity. Note: the value 2147483647 indicates Value\_Exceeds\_Maximum and the value -2147483648 indicates Value\_Exceeds\_Minimum.
- The type is int.
- The unit is faults/second.

#### Memory Pages Input/sec

- The number of pages read from the disk to resolve memory references to pages that were not in memory at the time of the reference. This counter includes paging traffic on behalf of the system cache to access file data for applications. This is an important counter to observe if you are concerned about excessive memory pressure (that is, thrashing), and the excessive paging that may result.
- The type is int.
- The unit is pages/second.

#### Memory Pool Paged Allocs

- The number of calls to allocate space in the system paged pool. Paged pool is a system memory area where space is acquired by operating system components as they accomplish their appointed tasks. Paged pool pages can be paged out to the paging file when not accessed by the system for sustained periods of time.
- The type is int.
- The unit is calls.

#### Memory Cache Faults/sec

- The average number of cache faults that have occurred on a system per second. Cache faults occur whenever the cache manager does not find a file's page in the immediate cache and must ask the memory manager to locate the page elsewhere in memory or on the disk so that it can be loaded in to the immediate cache.
- The type is int.
- The unit is faults/second.

#### Memory Working Set Total Bytes

- The number of bytes of real memory allocated to currently running processes.
- The type is double.
- The unit is bytes.

#### Memory Cache kBytes Peak

- The maximum number of KBs the system cache has used since startup.
- The type is double.
- The unit is kilobytes.

#### Memory Write Copies/sec

- The number of page faults that have been satisfied by making a copy of a page when an attempt to write to the page is made. This is an economical way of sharing data since the copy of the page is only made on an attempt to write to the page; otherwise, the page is shared.
- The type is int.
- The unit is faults/second.

#### Memory Working Set Total Usage Percentage

- The percent of bytes of real memory allocated to currently running processes. Calculated as  $100 * (\text{Total Working Set Bytes} / \text{Total Memory Bytes})$ .
- The type is int.
- The unit is percent.

#### Memory System Cache Resident Bytes

- The number of bytes from the system cache that are resident in physical memory. This does not include virtual memory not currently resident.
- The type is double.
- The unit is bytes.

#### Memory Commit Limit (Kilobytes)

- The number of KBs of virtual memory that can be committed on a system without extending the paging files.
- The type is double.
- The unit is kilobytes.

#### Memory Pool Nonpaged Bytes

- The number of bytes in the nonpaged pool, a system memory area where space is acquired by operating system components as they accomplish their appointed tasks. Nonpaged pool pages cannot be paged out to the paging file, but instead remain in main memory as long as they are allocated.
- The type is double.
- The unit is bytes.

#### Memory Free System Page Table Entries

- The number of page table entries a system is not currently using.
- The type is int.
- The unit is pageTables.

### **Component: ICMP Statistics**

Information about ICMP (Internet Control Message Protocol) messages that are used to convey the results of network commands, such as the PING command.

#### **Dimensions**

##### ICMP Parameter

- A special one for the front end to use as a column header.
- The type is string.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### node

- The managed system name.
- The type is string.

##### ICMP Value

- A special one for the front end to use as a column header.
- The type is string.

#### **Metrics**

##### ICMP Messages Sent/sec

- The rate that ICMP messages are attempted to be sent by the entity. The rate includes those messages sent in error.

- The type is int.
- The unit is messages/second.

#### ICMP Sent Echo Reply/sec

- The rate of ICMP Echo Reply messages sent.
- The type is int.
- The unit is messages/second.

#### ICMP Received Address Mask Reply

- The number of ICMP Address Mask Reply messages received.
- The type is int.
- The unit is messages.

#### ICMP Received Parameter Problem

- The number of ICMP Parameter Problem messages received.
- The type is int.
- The unit is messages/second.

#### ICMP Received Timestamp Reply/sec

- The rate of ICMP Timestamp Reply messages received.
- The type is int.
- The unit is messages/second.

#### ICMP Message Rate

- The total rate that ICMP messages that are received and sent by the entity. The rate includes those messages received or sent in error.
- The type is int.
- The unit is messages/second.

#### ICMP Sent Echo/sec

- The rate of ICMP Echo messages sent.
- The type is int.
- The unit is messages/second.

#### ICMP Received Timestamp/sec

- The rate of ICMP Timestamp (request) messages received.
- The type is int.
- The unit is messages/second.

#### ICMP Sent Timestamp/sec

- The rate of ICMP Timestamp (request) messages sent.
- The type is int.
- The unit is messages/second.

#### ICMP Messages Received Errors

- The number of ICMP messages that the entity received but determined as having errors (bad ICMP checksums, bad length, and so on).

- The type is int.
- The unit is errors.

#### ICMP Sent Source Quench

- The number of ICMP Source Quench messages sent.
- The type is int.
- The unit is messages.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### ICMP Received Echo Reply/sec

- The rate of ICMP Echo Reply messages received.
- The type is int.
- The unit is messages/second.

#### ICMP Received Redirect/sec

- The rate of ICMP Redirect messages received.
- The type is int.
- The unit is messages/second.

#### ICMP Sent Time Exceeded

- The number of ICMP Time Exceeded messages sent.
- The type is int.
- The unit is messages.

#### ICMP Sent Address Mask Reply

- The number of ICMP Address Mask Reply messages sent.
- The type is int.
- The unit is messages.

#### ICMP Sent Parameter Problem

- The number of ICMP Parameter Problem messages sent.
- The type is int.
- The unit is messages.

#### ICMP Received Time Exceeded

- The number of ICMP Time Exceeded messages received.
- The type is int.
- The unit is messages.

#### ICMP Received Destination Unreachable

- The number of ICMP Destination Unreachable messages received.

- The type is int.
- The unit is messages.

#### ICMP Sent Destination Unreachable

- The number of ICMP Destination Unreachable messages sent.
- The type is int.
- The unit is messages.

#### ICMP Sent Address Mask

- The number of ICMP Address Mask Request message sent.
- The type is int.
- The unit is messages.

#### ICMP Sent Redirect/sec

- The rate of ICMP Redirect messages sent.
- The type is int.
- The unit is messages/second.

#### ICMP Messages Outbound Errors

- The number of ICMP messages that this entity did not send due to problems discovered within ICMP, such as lack of buffers. This value must not include errors discovered outside the ICMP layer, such as the inability of IP to rout the resultant datagram. In some implementations, there might not be any types of error that contribute to the value of the counter.
- The type is int.
- The unit is errors.

#### ICMP Received Echo/sec

- The rate of ICMP Echo messages received.
- The type is int.
- The unit is messages/second.

#### ICMP Sent Timestamp Reply/sec

- The rate of ICMP Timestamp Reply messages sent.
- The type is int.
- The unit is messages/second.

#### ICMP Messages Received/sec

- The rate that ICMP messages are received by the entity. The rate includes those messages received in error.
- The type is int.
- The unit is messages/second.

#### ICMP Received Address Mask

- The number of ICMP Address Mask Request messages received.
- The type is int.
- The unit is messages.

#### ICMP Received Source Quench



- The number of ICMP Source Quench messages received.
- The type is int.
- The unit is messages.

### **Component: UDP Statistics**

Information about the datagram traffic for data using the UDP protocol.

#### **Dimensions**

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

UDP Value

- A special one for the front end to use as a column header.
- The type is string.

UDP Parameter

- A special one for the front end to use as a column header.
- The type is string.

node

- The managed system name.
- The type is string.

#### **Metrics**

UDP Datagrams Received/sec

- The rate that UDP datagrams are delivered to UDP users.
- The type is int.
- The unit is datagrams/second.

Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

UDP Datagrams No Port/sec

- The rate of received UDP datagrams for which there was no application at the destination port.
- The type is int.
- The unit is datagrams/second.

UDP Datagrams Received Errors

- The number of received UDP Datagrams that could not be delivered for reasons other than the lack of an application at the destination port.
- The type is int.
- The unit is datagrams.

#### UDP Datagram Rate

- The rate that UDP datagrams are sent or received by the entity.
- The type is int.
- The unit is datagrams/second.

#### UDP Datagrams Sent/sec

- The rate that UDP datagrams are sent from the entity.
- The type is int.
- The unit is datagrams/second.

### **Component: DNS Query**

Information about the DNS (Domain Name Server) server activity and performance.

#### **Dimensions**

##### DNS Query node

- The managed system name.
- The type is string.

##### DNS Query DNS Parameter

- A special one for the front end to use as a column header.
- The type is string.

##### DNS Query DNS Value

- A special one for the front end to use as a column header.
- The type is string.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### **Metrics**

##### DNS Query DNS Total Response Sent/sec

- The average number of responses sent by DNS server in each second.
- The type is int.
- The unit is responses/second.

##### DNS Query DNS Recursive Send TimeOuts

- The total number of recursive query sending timeouts.
- The type is int.
- The unit is queries.

##### DNS Query DNS Total Query Received/sec

- The average number of queries received by DNS server in each second.
- The type is int.
- The unit is queries/second.

##### DNS Query DNS TCP Query Received

- The total number of TCP queries received by DNS server.
- The type is int.
- The unit is queries.

#### DNS Query DNS Total Query Received

- The total number of queries received by DNS server.
- The type is int.
- The unit is queries.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### DNS Query DNS UDP Response Sent/sec

- The average number of UDP responses sent by DNS server in each second.
- The type is int.
- The unit is responses/second.

#### DNS Query DNS Recursive Query Failure

- The total number of recursive query failures.
- The type is int.
- The unit is queries.

#### DNS Query DNS UDP Query Received

- The total number of UDP queries received by DNS server.
- The type is int.
- The unit is queries.

#### DNS Query DNS TCP Response Sent

- The total number of TCP responses sent by DNS server.
- The type is int.
- The unit is responses.

#### DNS Query DNS Total Response Sent

- The total number of responses sent by DNS server.
- The type is int.
- The unit is responses.

#### DNS Query DNS UDP Query Received/sec

- The average number of UDP queries received by DNS server in each second.
- The type is int.
- The unit is queries/second.

#### DNS Query DNS Recursive Query Failure/sec

- The average number of recursive query failures in each second.
- The type is int.
- The unit is queries/second.

#### DNS Query DNS UDP Response Sent

- The total number of UDP responses sent by DNS server.
- The type is int.
- The unit is responses.

#### DNS Query DNS Recursive Queries/sec

- The average number of recursive queries received by DNS server in each second.
- The type is int.
- The unit is queries/second.

#### DNS Query DNS Recursive TimeOut/sec

- The average number of recursive query sending timeouts in each second.
- The type is int.
- The unit is queries/second.

#### DNS Query DNS TCP Query Received/sec

- The average number of TCP queries received by DNS server in each second.
- The type is int.
- The unit is queries/second.

#### DNS Query DNS Recursive Queries

- The total number of recursive queries received by DNS server.
- The type is int.
- The unit is queries.

#### DNS Query DNS TCP Response Sent/sec

- The average number of TCP responses sent by DNS server in each second.
- The type is int.
- The unit is responses/second.

### **Component: IIS Statistics**

Information about IIS memory usage and connection data.

#### **Dimensions**

node

- The managed system name.
- The type is string.

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

IIS Parameter

- A special one for the front end to use as a column header.
- The type is string.

#### IIS Value

- A special one for the front end to use as a column header.
- The type is string.

#### IIS Configured Size of Shared HTTP / FTP and Gopher Cache

- The configured maximum size of the shared HTTP, FTP, and Gopher memory cache.
- The type is int.

### **Metrics**

#### IIS Cache Misses

- The total number of times a file open, a directory listing, or a service-specific object request was not found in the cache.
- The type is int.
- The unit is count.

#### IIS File Expirations

- The number of times a portion of the memory cache has expired due to the file or directory changes in an IIS directory tree. Note that this attribute is not available on systems running Windows 2000 with IIS 5.0.
- The type is int.
- The unit is flushes.

#### IIS Measured Async I/O Bandwidth Usage

- The number of measured bandwidth of asynchronous I/O averaged over a minute.
- The type is int.
- The unit is bandwidth/minute.

#### IIS Total Allowed Async Requests

- The total number of asynchronous I/O requests allowed by the bandwidth throttler.
- The type is int.
- The unit is requests.

#### IIS Total Blocked Async I/O Requests

- The total number of asynchronous I/O requests blocked by the bandwidth throttler.
- The type is int.
- The unit is requests.

#### IIS Cache Objects

- The number of objects cached by all of the IIS. The objects include file handle tracking objects, directory listing objects, and service-specific objects.
- The type is int.
- The unit is objects.

#### IIS Cache Used

- The total number of bytes currently containing cached data in the shared memory cache. This includes directory listings, file handle tracking, and service-specific objects.
- The type is int.
- The unit is bytes.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### IIS Total Rejected Async Requests

- The total number of asynchronous I/O requests rejected by the bandwidth throttler.
- The type is int.
- The unit is requests.

#### IIS Cached File Handles

- The number of open file handles cached by all of the IIS. Note that this attribute is not available on systems running Windows 2000 with IIS 5.0.
- The type is int.
- The unit is handles.

#### IIS Statistics Cache Hits

- The total number of times a file open, a directory listing, or a service-specific object request was found in the cache.
- The type is int.
- The unit is count.

#### IIS Cached Directory Listings

- The number of directory listing objects cached by all of the IIS. Note that this attribute is not available on systems running Windows 2000 with IIS 5.0.
- The type is int.
- The unit is listings.

#### IIS Statistics Cache Hits %

- The ratio of cache hits to all cache requests.
- The type is int.
- The unit is percent.

#### IIS Current Blocked Async I/O Requests

- The current number of asynchronous I/O requests blocked by the bandwidth throttler.
- The type is int.
- The unit is requests.

### **Component: Active Tasks running on VCenter Server**

Use the Event Log data set to create situations about actual records that are written to any Windows Event logs, such as date and time of the event and event identification information. Event Log is a

multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

### **Dimensions**

#### Event Log User Name

- The name of the user whose information you are monitoring.
- The type is string.

#### Event Log Event Time

- The time for when the event was logged.
- The type is string.

#### Event Log Event ID

- The identification code of the event you are monitoring.
- The type is int.

#### Event Log Event Time Stamp

- The date and time the event you are monitoring is logged.
- The type is timestamp.

#### Event Log Source (Unicode)

- The software that logged the event, which can be an application name or a component of the system in UTF8.
- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### Event Log Event Description

- A description of the event you are monitoring. Note that this attribute displays only ANSI strings. Description (Unicode) serves as a Unicode version of this attribute.
- The type is string.

#### Event Log Event Class

- The classification of the event as defined by the source.
- The type is string.

#### Event Log Event Severity

- The severity level of the event you are monitoring.
- The type is string.

#### Event Log Computer Name

- The name of the computer where the event occurs.
- The type is string.

#### Event Log User (Unicode)

- The user name in UTF8. Valid format is a text string of up to 52 bytes.

- The type is string.

#### Event Log Record Number

- The identifier for the event within the Windows NT event log file (specific to the log file). This attribute is used together with the log file name to uniquely identify an instance of this class.
- The type is double.

#### Event Log Log Name

- The name of a log. Valid format is a text string of up to 32 characters. The log names are case sensitive. Application is an example of a valid log name.
- The type is string.

#### Event Log Component/Application Name

- The name of the application or component that logged the event you are monitoring.
- The type is string.

#### Event Log Category (Unicode)

- The classification of the event as defined by the source in UTF8.
- The type is string.

#### Event Log Description (Unicode)

- A description of the event you are monitoring in UTF8.
- The type is string.

#### Event Log Event Date

- The date for when the event was logged.
- The type is timestamp.

#### node

- The managed system name.
- The type is string.

#### Event Log Log Name (Unicode)

- The Log Name in UTF8. Valid format is a text string of up to 392 bytes. The log names are case sensitive. Application is an example of a valid log name.
- The type is string. This is a key dimension.

#### Event Log Event ID (String)

- Event ID represented as a string.
- The type is string.

### Metrics

#### Event Log Duplicate Record Count

- The number of duplicate records in the NT Event Log. Note: the value 9223372036854775807 indicates Value\_Exceeds\_Maximum and the value -9223372036854775808 indicates Value\_Exceeds\_Minimum. Controlled by these agent environment settings:
  - NT\_LOG\_THROTTLE
  - NT\_{ *Event Log Name* }\_LOG\_THROTTLE



- NT\_APPLICATION\_LOG\_THROTTLE
- NT\_SYSTEM\_LOG\_THROTTLE
- NT\_SECURITY\_LOG\_THROTTLE
- NT\_DNS\_LOG\_THROTTLE
- NT\_DIRSERVICE\_LOG\_THROTTLE
- NT\_FILEREPSRV\_LOG\_THROTTLE

This field does not indicate the number of duplicate record counts since the agent started. Its range is limited to a situation that is a pure event. In case of a storm in the event logs, the attribute might assume a value greater than 0. The value of the Duplicate Record Count is aleatory as it does not refer to the past events in the event log, but instead, to the sum of events passed by the system each time that the agent awakes. For example, the following situation: (EVENTID EQ 4625) AND (DUPCNT > 2) does not trigger when the number of event IDs that equal 4625 are greater than 2 since the agent started, but only when 2 events with IDs equal to 4625 occur within the same set of events (a rare circumstance).

- The type is double.
- The unit is duplicates.

Row Number

- Row Number.
- The type is int.
- The unit is row.

### **Component: Active Tasks running on VCenter Server**

Use the Job Object data set to create situations that monitor job kernel objects, the system resources a job consumes, and the number of processes a job contains. Job Object is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

Job Object Process

- Name of process in UTF8. The maximum process name size is defined by MAX\_PATH.
- The type is string.

Job Object Priority Base

- The current base priority of this process. Threads within a process can raise and lower their own base priority relative to the process' base priority. Valid values are positive integers in the range 0 to 2147483647 and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions.
- The type is int.

Job Object Job Kernel Object

- Name of Job kernel object. Valid format is a text string of up to 64 characters.
- The type is string. This is a key dimension.

Job Object Name (Unicode)

- Instance name (Job Object) in UTF8. The maximum job object name size is defined by MAX\_PATH.
- The type is string. This is a key dimension.

Time Stamp

- The date and time the agent collects information as set on the monitored system.

- The type is timestamp.

node

- The managed system name. The form should be *hostname : agent\_code* . Examples include spark:KNT or deux.raleigh.ibm.com:KNT.
- The type is string.

Job Object Creating Process ID

- The Process ID of the creating process. Note that the creating process might have terminated since this process was created, and so this value might no longer identify a running process.
- The type is double.

node

- The managed system name.
- The type is string.

Job Object ID Process

- The unique identifier of this process. ID Process numbers are reused, so they only identify a process for the lifetime of that process.
- The type is double. This is a key dimension.

## Metrics

Job Object IO Data Bytes/sec

- The rate the process is reading and writing bytes in I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/O's.
- The type is double.
- The unit is bytes/second.

Job Object Virtual Bytes

- The current size in bytes of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and by using too much, the process can limit its ability to load libraries.
- The type is double.
- The unit is bytes.

Job Object Total mSec Kernel Mode

- Shows the number of milliseconds of kernel mode processor time used by all the processes in the Job object, including those that have terminated or that are no longer associated with the Job object, since the Job object was created. This attribute is the 64-bit version of Total mSec Kernel Mode.
- The type is double.
- The unit is milliseconds.

Job Object IO Write Operations/sec

- The rate the process is issuing write I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/O's.
- The type is double.
- The unit is operations/second.

#### Job Object Page File Bytes

- The current number of bytes this process has used in the paging file(s). Paging files are used to store pages of memory used by the process that are not contained in other files. Paging files are shared by all processes, and lack of space in paging files can prevent other processes from allocating memory.
- The type is double.
- The unit is bytes.

#### Job Object Virtual Bytes Peak

- Virtual Bytes Peak is the maximum number of bytes of virtual address space the process has used at any one time. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is however finite, and by using too much, the process might limit its ability to load libraries.
- The type is double.
- The unit is bytes.

#### Job Object Process Count Active

- Shows the number of processes that are currently associated with the Job object.
- The type is int.
- The unit is processes.

#### Job Object Current % Kernel Mode Time

- Shows the percentage of the monitoring interval that the processes in the Job object spent executing code in kernel or privileged mode. Valid values are positive integers in the range 0 to 2147483647 and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions.
- The type is int.
- The unit is percent.

#### Job Object Total mSec User Mode

- Shows the number of milliseconds of user mode processor time used by all the processes in the Job object, including those that have terminated or that are no longer associated with the Job object, since the Job object was created. This attribute is the 64-bit version of Total\_mSec\_User\_Mode.
- The type is double.
- The unit is milliseconds.

#### Job Object Virtual kBytes Peak

- Virtual Bytes Peak in KBs.
- The type is double.
- The unit is kilobytes.

#### Job Object Thread Count

- The number of threads currently active in this process. An instruction is the basic unit of execution in a processor, and a thread is the object that executes instructions. Every running process has at least one thread.
- The type is int.
- The unit is threads.

#### This Period mSec Processor (Superseded)

- Shows the number of milliseconds of processor time used by all the processes in the Job object, including those that have terminated or that are no longer associated with the Job object, since the last time a time limit on the Job was established.
- The type is int.
- The unit is milliseconds.

#### Job Object This Period mSec Processor

- Shows the number of milliseconds of processor time used by all the processes in the Job object, including those that have terminated or that are no longer associated with the Job object, since the last time a time limit on the Job was established. This attribute is the 64-bit version of This Period mSec Processor.
- The type is double.
- The unit is milliseconds.

#### Job Object Page File kBytes Peak

- Page File Bytes Peak in KBs.
- The type is double.
- The unit is kilobytes.

#### Job Object Current % User Mode Time

- Shows the percentage of the monitoring interval that the processes in the Job object spent executing code in user mode.
- The type is int.
- The unit is percent.

#### Job Object Page Faults/sec

- The rate Page Faults occur in the threads executing in this process. A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory. This does not cause the page to be fetched from disk if it is on the standby list and hence already in main memory, or if it is in use by another process with whom the page is shared.
- The type is int.
- The unit is faults/second.

#### Job Object Working Set Peak

- The maximum number of bytes in the Working Set of this process at any point in time. The Working Set is the set of memory pages touched recently by the threads in the process. If free memory in the computer is above a threshold, pages are left in the Working Set of a process even if they are not in use. When free memory falls below a threshold, pages are trimmed from Working Sets. If they are needed, they are soft-faulted back in to the Working Set before they leave main memory.
- The type is double.
- The unit is bytes.

#### Job Object Handle Count

- The total number of handles currently open by this process. This number is the sum of the handles currently open by each thread in this process.
- The type is int.
- The unit is handles.

#### Job Object Virtual kBytes

- Virtual Bytes in KBs.
- The type is double.
- The unit is kilobytes.

#### Job Object Working Set

- The current number of bytes in the Working Set of this process. The Working Set is the set of memory pages touched recently by the threads in the process. If free memory in the computer is above a threshold, pages are left in the Working Set of a process even if they are not in use. When free memory falls below a threshold, pages are trimmed from Working Sets. If they are needed, they are soft-faulted back in to the Working Set before they leave main memory.
- The type is double.
- The unit is bytes.

#### Job Object Private Bytes

- The current number of bytes this process has allocated that cannot be shared with other processes.
- The type is double.
- The unit is bytes.

#### Job Object % Processor Time

- The percentage of elapsed time that all of the threads of this process used the processor to execute instructions. An instruction is the basic unit of execution in a computer, a thread is the object that executes instructions, and a process is the object created when a program is run. Code executed to handle some hardware interrupts and trap conditions are included in this count. On Multi-processor systems the maximum value of the counter is 100 % times the number of processors.
- The type is int.
- The unit is percent.

#### Job Object % User Time

- The percentage of elapsed time that this threads of this process have spent executing code in user mode. Applications, environment subsystems and integral subsystems execute in user mode. Code executing in user mode cannot damage the integrity of the Windows Server Executive, Kernel, and device drivers. Unlike some early operating systems, Windows NT and higher versions of Windows Servers use process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Windows Server on behalf of your application might appear in other subsystem processes in addition to the privileged time in your process.
- The type is int.
- The unit is percent.

#### Job Object This Period mSec Kernel Mode (Superseded)

- Shows the number of milliseconds of kernel mode processor time used by all the processes in the Job object, including those that have terminated or that are no longer associated with the Job object, since the last time a time limit on the Job was established.
- The type is int.
- The unit is milliseconds.

#### Job Object This Period mSec User Mode (Superseded)

- Shows the number of milliseconds of user mode processor time used by all the processes in the Job object, including those that have terminated or that are no longer associated with the Job object, since the last time a time limit on the Job was established.
- The type is int.
- The unit is milliseconds.

#### Job Object Pool Nonpaged Bytes

- The number of bytes in the nonpaged pool, an area of system memory (physical memory used by the operating system) for objects that cannot be written to disk, but must remain in physical memory as long as they are allocated.
- The type is int.
- The unit is bytes.

#### Job Object Total mSec Processor

- Shows the number of milliseconds of processor time used by all the processes in the Job object, including those that have terminated or that are no longer associated with the Job object, since the Job object was created. This attribute is the 64-bit version of Total mSec Processor.
- The type is double.
- The unit is milliseconds.

#### Job Object IO Read Operations/sec

- The rate the process is issuing read I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/O's.
- The type is double.
- The unit is operations/second.

#### Job Object Pool Paged Bytes

- The number of bytes in the paged pool, an area of system memory (physical memory used by the operating system) for objects that can be written to disk when they are not being used.
- The type is int.
- The unit is bytes.

#### Job Object IO Write Bytes/sec

- The rate the process is writing bytes to I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/O's.
- The type is double.
- The unit is bytes/second.

#### Job Object Pages/sec

- Shows the page fault rate of all the processes in the Job object.
- The type is int.
- The unit is pages/second.

#### Job Object Current % Processor Time

- Shows the percentage of the monitoring interval that the process in the Job object spent executing code.
- The type is int.
- The unit is percent.

#### Job Object Total mSec User Mode (Superseded)

- Shows the number of milliseconds of user mode processor time used by all the processes in the Job object, including those that have terminated or that are no longer associated with the Job object, since the Job object was created.
- The type is int.
- The unit is milliseconds.

#### Job Object IO Other Operations/sec

- The rate the process is issuing I/O operations that are neither a read or a write operation. An example of this type of operation would be a control function. This counter counts all I/O activity generated by the process to include file, network and device I/O's.
- The type is double.
- The unit is operations/second.

#### Job Object Elapsed Time

- The total elapsed time (in seconds) this process has been running.
- The type is double.
- The unit is seconds.

#### Job Object Private kBytes

- Private Bytes in KBs.
- The type is double.
- The unit is kilobytes.

#### Job Object This Period mSec User Mode

- Shows the number of milliseconds of user mode processor time used by all the processes in the Job object, including those that have terminated or that are no longer associated with the Job object, since the last time a time limit on the Job was established. This attribute is the 64-bit version of This Period mSec User Mode.
- The type is double.
- The unit is milliseconds.

#### Job Object Page File kBytes

- Page File Bytes in KBs.
- The type is double.
- The unit is kilobytes.

#### Job Object IO Other Bytes/sec

- The rate the process is issuing bytes to I/O operations that do not involve data such as control operations. This counter counts all I/O activity generated by the process to include file, network and device I/O's.
- The type is double.
- The unit is bytes/second.

#### Job Object This Period mSec Kernel Mode

- Shows the number of milliseconds of kernel mode processor time used by all the processes in the Job object, including those that have terminated or that are no longer associated with the Job object, since the last time a time limit on the Job was established. This attribute is the 64-bit version of This Period mSec Kernel Mode.

- The type is double.
- The unit is milliseconds.

#### Job Object IO Data Operations/sec

- The rate the process is issuing read and write I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/O's.
- The type is double.
- The unit is operations/second.

#### Job Object Process Count Total

- Shows the number of processes, both active and terminated, that are or have been associated with the Job object.
- The type is int.
- The unit is processes.

#### Total mSec Processor (Superseded)

- Shows the number of milliseconds of processor time used by all the processes in the Job object, including those that have terminated or that are no longer associated with the Job object, since the Job object was created.
- The type is int.
- The unit is milliseconds.

#### Job Object % Privileged Time

- The percentage of elapsed time that the threads of the process have spent executing code in privileged mode. When a Windows Server's service is called, the service often runs in Privileged Mode to gain access to system-private data. Such data is protected from access by threads executing in user Mode. Calls to the system can be explicit or implicit, such as page faults or interrupts. Unlike some early operating systems, Windows NT and higher versions of Windows Servers use process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Windows Server on behalf of your application might appear in other subsystem processes in addition to the privileged time in your process.
- The type is int.
- The unit is percent.

#### Job Object Total mSec Kernel Mode (Superseded)

- Shows the number of milliseconds of kernel mode processor time used by all the processes in the Job object, including those that have terminated or that are no longer associated with the Job object, since the Job object was created.
- The type is int.
- The unit is milliseconds.

#### Job Object Page File Bytes Peak

- The maximum number of bytes this process has used in the paging file(s). Paging files are used to store pages of memory used by the process that are not contained in other files. Paging files are shared by all processes, and lack of space in paging files can prevent other processes from allocating memory.
- The type is double.
- The unit is bytes.



#### Job Object IO Read Bytes/sec

- The rate the process is reading bytes from I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/O's.
- The type is double.
- The unit is bytes/second.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### Job Object Process Count Terminated

- Shows the number of processes that have been terminated because of a limit violation.
- The type is int.
- The unit is processes.

### **Component: Active Tasks running on VCenter Server**

Use the Registry data set to monitor for specific values or changes in registry data.

#### **Dimensions**

##### node

- The managed system name.
- The type is string.

##### Registry Numeric Value

- The Registry Numeric Data Value. Contains numeric data for any registry entry whose type is defined as any of the numeric data types. These include: Binary, DWORD, Big Endian, QWORD.
- The type is int.

##### Registry String Value

- The Registry String Data Value. Contains string data for any registry entry whose type is defined as any of the string data types. These include: String, Expandable String, Multiple String.
- The type is string.

##### Registry Type

- The Registry value type.
- The type is int.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### Registry Path Name

- The Path is the concatenation of the registry sub-key and the registry value name.
- The type is string. This is a key dimension.

##### Registry Root Key Name

- The Registry Root key name. Valid values are positive integers representing any of the enumerated values for Windows Registry keys. When creating a situation using this attribute field, if you are using a remote machine, then only HKEY\_LOCAL\_MACHINE and HKEY\_USER are allowed.
- The type is int.

### Metrics

Row Number

- Row Number.
- The type is int.
- The unit is row.

### Component: VM Memory

Information about the memory statistics for this virtual machine.

#### Dimensions

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

node

- The managed system name.
- The type is string.

#### Metrics

VM Memory Memory Overhead in MB

- The amount of overhead memory associated with this virtual machine consumed on the host system.
- The type is int.
- The unit is megabytes.

VM Memory Memory Shared in MB

- The amount of physical memory associated with this virtual machine that is copy-on-write (COW) shared on the host.
- The type is int.
- The unit is megabytes.

VM Memory Memory Swapped in MB

- The amount of memory associated with this virtual machine that has been swapped by ESX.
- The type is int.
- The unit is megabytes.

VM Memory Memory Reservation in MB

- The minimum amount of memory that is guaranteed to the virtual machine. Assigning a Memory Reservation ensures that even as other virtual machines on the same host consume memory, there is still a certain minimum amount for this virtual machine.
- The type is int.
- The unit is megabytes.

#### VM Memory Memory Mapped in MB

- The mapped memory size of this virtual machine. This is the current total amount of guest memory that is backed by physical memory. This number might include pages of memory shared between multiple virtual machines and thus might overestimate the amount of physical host memory consumed by this virtual machine.
- The type is int.
- The unit is megabytes.

#### VM Memory Memory Ballooned in MB

- The amount of memory that has been reclaimed from this virtual machine via the VMware Memory Balloon mechanism.
- The type is int.
- The unit is megabytes.

#### VM Memory Memory Active in MB

- The estimated amount of memory that the virtual machine is actively using.
- The type is int.
- The unit is megabytes.

#### VM Memory Memory Shared Saved in MB

- The estimated amount of physical memory on the host saved from copy-on-write (COW) shared guest physical memory.
- The type is int.
- The unit is megabytes.

#### VM Memory Memory Limit in MB

- The maximum amount of memory that is allowed to the virtual machine. Assigning a Memory Limit ensures that this virtual machine never consumes more than a certain amount of the allowed memory. By limiting the amount of memory consumed, a portion of this shared resource is allowed to other virtual machines.
- The type is int.
- The unit is megabytes.

#### VM Memory Memory Shares

- The number of memory shares allocated to the virtual machine.
- The type is int.
- The unit is shares.

#### VM Memory Memory Target Size

- The memory target size.
- The type is int.
- The unit is megabytes.

#### VM Memory Memory Used in MB

- The estimated amount of physical host memory currently consumed for this virtual machine's physical memory. This estimation is the same as (Memory Mapped in MB) - (Memory Shared Saved in MB).
- The type is int.

- The unit is megabytes.

## **Component: Redirector**

Information about IO Statistics of Processes.

### **Dimensions**

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

Redirector Parameter

- A special one for the front end to use as a column header.
- The type is string.

node

- The managed system name.
- The type is string.

Redirector Value

- A special one for the front end to use as a column header.
- The type is string.

### **Metrics**

Redirector Server Reconnects

- The number of times your redirector has had to reconnect to a server in order to complete a new active request. You can be disconnected by the server if you remain inactive for too long. Locally even if all your remote files are closed, the redirector keeps your connections intact for (nominally) 10 minutes. Such inactive connections are called Dormant Connections. Reconnecting is expensive in time.
- The type is int.
- The unit is reconnects.

Redirector Write Bytes Network/sec

- The rate at which applications are writing data across the network. This occurs when the file system cache is bypassed, such as for Named Pipes or Transactions, or when the cache writes the bytes to disk to make room for other data. Dividing this counter by Bytes Transmitted/sec indicates the proportion of application data being to the network (see Transmitted Bytes/sec). This attribute is the 64-bit version of Write Bytes Network/Sec.
- The type is double.
- The unit is bytes/second.

Redirector Write Bytes Network/sec (Superseded)

- The rate at which applications are writing data across the network. This occurs when the file system cache is bypassed, such as for Named Pipes or Transactions, or when the cache writes the bytes to disk to make room for other data. Dividing this counter by Bytes Transmitted/sec indicates the proportion of application data being to the network (see Transmitted Bytes/sec).
- The type is int.
- The unit is bytes/second.

Redirector Read Packets/sec

- The rate at which read packets are being placed on the network. Each time a single packet is sent with a request to read data remotely, this counter is incremented by one.
- The type is int.
- The unit is packets/second.

#### Redirector Read Bytes Network/sec (Superseded)

- The rate at which applications are reading data across the network. This occurs when data sought in the file system cache is not found there and must be retrieved from the network. Dividing this value by Bytes Received/sec indicates the proportion of application data traveling across the network (see Bytes Received/sec).
- The type is int.
- The unit is bytes/second.

#### Redirector Bytes Received/sec

- The rate of bytes coming in to the redirector from the network. It includes all application data as well as network protocol information (such as packet headers). This attribute is the 64-bit version of Bytes Received/Sec.
- The type is double.
- The unit is bytes/second.

#### Redirector Packets/sec (Superseded)

- The rate the redirector is processing data packets. One packet includes many bytes, but each packet has protocol overhead. You can determine the efficiency of this path by dividing the Bytes/sec by this counter to determine the average number of bytes transferred/packet. You can also divide this counter by Operations/sec to determine the average number of packets per operation, another measure of efficiency.
- The type is int.
- The unit is packets/second.

#### Redirector Read Bytes Non-Paging/sec (Superseded)

- The bytes read by the redirector in response to normal file requests by an application when they are redirected to come from another computer. In addition to file requests, this counter includes other methods of reading across the network such as Named Pipes and Transactions. This counter does not count network protocol information, just application data.
- The type is int.
- The unit is bytes/second.

#### Redirector Read Operations Random/sec

- The rate at which, on a file-by-file basis, reads are made that are not sequential. If a read is made using a particular file handle, and then is followed by another read that is not immediately the contiguous next byte, this counter is incremented by one.
- The type is int.
- The unit is operations/second.

#### Redirector Packets Transmitted/sec

- The rate at which the redirector is sending packets (also called SMBs or Server Message Blocks). Network transmissions are divided in to packets. The average number of bytes transmitted in a packet can be obtained by dividing Bytes Transmitted/sec by this counter. This attribute is the 64-bit version of Packets\_Transmitted/Sec.
- The type is double.

- The unit is packets/second.

#### Redirector File Read Operations/sec

- The rate at which applications are asking the redirector for data. Each call to a file system or similar Application Program Interface (API) call counts as one operation.
- The type is int.
- The unit is operations/second.

#### Redirector Connects Windows NT

- Counts the connections to Windows 2000 or earlier computers.
- The type is int.
- The unit is connections.

#### Redirector Connects LAN Manager 2.0

- Counts connections to LAN Manager 2(dot)0 servers, including LMX servers.
- The type is int.
- The unit is connections.

#### Redirector File Write Operations/sec

- The rate at which applications are sending data to the redirector. Each call to a file system or similar Application Program Interface (API) call counts as one operation.
- The type is int.
- The unit is operations/second.

#### Redirector Packets Received/sec (Superseded)

- The rate at which the redirector is receiving packets (also called SMBs or Server Message Blocks). Network transmissions are divided into packets. The average number of bytes received in a packet can be obtained by dividing Bytes Received/sec by this counter. Some packets received might not contain incoming data (for example, an acknowledgment to a write made by the redirector counts as an incoming packet).
- The type is int.
- The unit is packets/second.

#### Redirector Bytes Total/sec

- The rate the redirector is processing data bytes. This includes all application and file data in addition to protocol information such as packet headers. This attribute is the 64-bit version of Bytes Total/Sec.
- The type is double.
- The unit is bytes/second.

#### Redirector Bytes Total/sec (Superseded)

- The rate the redirector is processing data bytes. This includes all application and file data in addition to protocol information such as packet headers.
- The type is int.
- The unit is bytes/second.

#### Redirector File Data Operations/sec

- The rate at which the redirector is processing data operations. One operation can include many bytes, since each operation has overhead. The efficiency of this path can be determined by

dividing the Bytes/sec by this counter to obtain the average number of bytes transferred per operation.

- The type is int.
- The unit is operations/second.

#### Redirector Connects LAN Manager 2.1

- Counts connections to LAN Manager 2(dot)1 servers, including LMX servers.
- The type is int.
- The unit is connections.

#### Redirector Network Errors/sec

- The rate at which serious unexpected errors are occurring. Such errors generally indicate that the redirector and one or more servers are having communication difficulties. For example, an SMB (Server Manager Block) protocol error is a Network Error. An entry is written to the System Event Log and provide details.
- The type is int.
- The unit is errors/second.

#### Redirector Read Bytes Network/sec

- The rate at which applications are reading data across the network. This occurs when data sought in the file system cache is not found there and must be retrieved from the network. Dividing this value by Bytes Received/sec indicates the proportion of application data traveling across the network (see Bytes Received/sec). This attribute is the 64-bit version of Read Bytes Network/Sec.
- The type is double.
- The unit is bytes/second.

#### Redirector Redirector Bytes Transmitted/sec

- The rate at which bytes are leaving the redirector to the network. It includes all application data as well as network protocol information (such as packet headers). This attribute is the 64-bit version of Bytes Transmitted/Sec.
- The type is double.
- The unit is bytes/second.

#### Redirector Writes Denied/sec

- The rate at which the server is unable to accommodate requests for Raw Writes. When a write is much larger than the server's negotiated buffer size, the redirector requests a raw write which, if granted, would permit the transfer of the data without lots of protocol overhead on each packet. To accomplish this the server must lock out other requests, so the request is denied if the server is really busy.
- The type is int.
- The unit is requests/second.

#### Redirector Server Sessions

- The total number of security objects the redirector has managed. For example, a logon to a server followed by a network access to the same server establishes one connection, but two sessions.
- The type is int.
- The unit is sessions.

#### Redirector Write Bytes Paging/sec

- The rate at which the redirector is attempting to write bytes changed in the pages being used by applications. The program data changed by modules (such as programs and libraries) that were loaded over the network are 'paged out' when no longer needed. Other output pages come from the file system cache (see Write Bytes Cache/sec). This attribute is the 64-bit version of Write Bytes Paging/Sec.
- The type is double.
- The unit is bytes/second.

#### Redirector Read Packets Small/sec

- The rate at which reads less than one-fourth of the server negotiated buffer size are made by applications. Too many of these could indicate a waste of buffers on the server. This counter is incremented once for each read. It does not count packets.
- The type is int.
- The unit is packets/second.

#### Redirector High % Bytes/Sec

- The percentage of network card bandwidth that is used by the redirector (workstation) service.
- The type is int.
- The unit is percent/second.

#### Redirector Write Packets Small/sec

- The rate at which writes are made by applications that are less than one-fourth of the negotiated buffer size of the server. Too many of these could indicate a waste of buffers on the server. This counter is incremented once for each write: it counts writes, not packets.
- The type is int.
- The unit is writes/second.

#### Redirector Current Commands

- Counts the number of requests to the redirector that are currently queued for service. If this number is much larger than the number of network adapter cards installed in the computer, then the networks and the servers being accessed are bottlenecked.
- The type is int.
- The unit is requests.

#### Redirector Read Bytes Cache/sec (Superseded)

- The rate at which applications are accessing the file system cache by using the redirector. Some of these data requests are satisfied by retrieving the data from the cache. Requests that miss the cache cause a page fault (see Read Bytes Paging/sec).
- The type is int.
- The unit is bytes/second.

#### Redirector Write Bytes Paging/sec (Superseded)

- The rate at which the redirector is attempting to write bytes changed in the pages being used by applications. The program data changed by modules (such as programs and libraries) that were loaded over the network are 'paged out' when no longer needed. Other output pages come from the file system cache (see Write Bytes Cache/sec). Valid values are positive integers in the range 0 to 2147483647 and can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions.
- The type is int.



- The unit is bytes/second.

#### Redirector Read Bytes Paging/sec

- The rate at which the redirector is attempting to read bytes in response to page faults. Page faults are caused by loading of modules (such as programs and libraries), by a miss in the cache (see Read Bytes Cache/sec), or by files directly mapped in to the address space of applications. This attribute is the 64-bit version of Read Bytes Paging/Sec.
- The type is double.
- The unit is bytes/second.

#### Redirector Packets Transmitted/sec (Superseded)

- The rate at which the redirector is sending packets (also called SMBs or Server Message Blocks). Network transmissions are divided in to packets. The average number of bytes transmitted in a packet can be obtained by dividing Bytes Transmitted/sec by this counter.
- The type is int.
- The unit is packets/second.

#### Redirector Write Bytes Non-Paging/sec (Superseded)

- The rate at which bytes are written by the redirector in response to normal file outputs by an application when they are redirected to another computer. In addition to file requests, this count includes other methods of writing across the network, such as Named Pipes and Transactions. This counter does not count network protocol information, just application data.
- The type is int.
- The unit is bytes/second.

#### Redirector Reads Denied/sec

- The rate at which the server is unable to accommodate requests for raw reads. When a read is much larger than the server's negotiated buffer size, the redirector requests a raw read which, if granted, would permit the transfer of the data without a lot of protocol overhead on each packet. To accomplish this, the server must lock out other requests, therefore the request is denied if the server is very busy.
- The type is int.
- The unit is requests/second.

#### Redirector Connects Core

- The number of connections you have to servers running the original MS-Net SMB protocol, including MS-Net itself, Xenix, and VAXs.
- The type is int.
- The unit is connections.

#### Redirector Read Bytes Paging/sec (Superseded)

- The rate at which the redirector is attempting to read bytes in response to page faults. Page faults are caused by loading of modules (such as programs and libraries), by a miss in the cache (see Read Bytes Cache/sec), or by files directly mapped in to the address space of applications.
- The type is int.
- The unit is bytes/second.

#### Redirector Server Disconnects

- The number of times a server has disconnected your redirector. (See also Server Reconnects.)
- The type is int.

- The unit is disconnects.

#### Redirector Read Bytes Non-Paging/sec

- The bytes read by the redirector in response to normal file requests by an application when they are redirected to come from another computer. In addition to file requests, this counter includes other methods of reading across the network such as Named Pipes and Transactions. This counter does not count network protocol information, just application data. This attribute is the 64-bit version of Read Bytes Non-Paging/Sec.
- The type is double.
- The unit is bytes/second.

#### Redirector Read Bytes Cache/sec

- The rate at which applications are accessing the file system cache by using the redirector. Some of these data requests are satisfied by retrieving the data from the cache. Requests that miss the cache cause a page fault (see Read Bytes Paging/sec). This attribute is the 64-bit version of Read Bytes Cache/Sec.
- The type is double.
- The unit is bytes/second.

#### Redirector Writes Large/sec

- The rate at which writes are made by applications that are more than twice the negotiated buffer size of the server. Too many of these could place a strain on server resources. This counter is incremented once for each write: it counts writes, not packets.
- The type is int.
- The unit is requests/second.

#### Redirector Server Sessions Hung

- The number of active sessions that are timed out and unable to proceed due to a lack of response from the remote server.
- The type is int.
- The unit is sessions.

#### Redirector Write Bytes Non-Paging/sec

- The rate at which bytes are written by the redirector in response to normal file outputs by an application when they are redirected to another computer. In addition to file requests, this count includes other methods of writing across the network, such as Named Pipes and Transactions. This counter does not count network protocol information, just application data. This attribute is the 64-bit version of Write Bytes Non-Paging/Sec.
- The type is double.
- The unit is bytes/second.

#### Redirector Packets/sec

- The rate the redirector is processing data packets. One packet includes many bytes, but each packet has protocol overhead. You can determine the efficiency of this path by dividing the Bytes/sec by this counter to determine the average number of bytes transferred/packet. You can also divide this counter by Operations/sec to determine the average number of packets per operation, another measure of efficiency. This attribute is the 64-bit version of Packets/Sec.
- The type is double.
- The unit is packets/second.

#### Redirector Bytes Received/sec (Superseded)

- The rate of bytes coming in to the redirector from the network. It includes all application data as well as network protocol information (such as packet headers).
- The type is int.
- The unit is bytes/second.

#### Redirector Bytes Transmitted/sec (Superseded)

- The rate at which bytes are leaving the redirector to the network. It includes all application data as well as network protocol information (such as packet headers).
- The type is int.
- The unit is bytes/second.

#### Redirector High Current Mod

- The average number of requests, per network interface card, to the redirector that are currently queued for service. This value is calculated as Current Commands NIC Count.
- The type is int.
- The unit is requests.

#### Redirector Write Operations Random/sec

- The rate at which, on a file-by-file basis, writes are made that are not sequential. If a write is made using a particular file handle, and then is followed by another write that is not immediately the next contiguous byte, this counter is incremented by one.
- The type is int.
- The unit is operations/second.

#### Redirector Packets Received/sec

- The rate at which the redirector is receiving packets (also called SMBs or Server Message Blocks). Network transmissions are divided in to packets. The average number of bytes received in a packet can be obtained by dividing Bytes Received/sec by this counter. Some packets received might not contain incoming data (for example, an acknowledgment to a write made by the redirector counts as an incoming packet). This attribute is the 64-bit version of Packets Received/Sec.
- The type is double.
- The unit is packets/second.

#### Redirector Write Bytes Cache/sec (Superseded)

- The rate at which applications on your computer are writing to the file system cache by using the Redirector. The data might not leave your computer immediately; it can be retained in the cache for further modification before being written to the network. This saves network traffic. Each write of a byte in to the cache is counted here.
- The type is int.
- The unit is bytes/second.

#### Redirector Write Bytes Cache/sec

- The rate at which applications on your computer are writing to the file system cache by using the redirector. The data might not leave your computer immediately; it can be retained in the cache for further modification before being written to the network. This saves network traffic. Each write of a byte in to the cache is counted here.
- The type is double.
- The unit is bytes/second.

#### Redirector Reads Large/sec

- The rate at which reads more than twice the negotiated buffer size of the server are made by applications. Too many of these could place a strain on server resources. This counter is incremented once for each read. It does not count packets.
- The type is int.
- The unit is requests/second.

#### Redirector Write Packets/sec

- The rate at which writes are being sent to the network. Each time a single packet is sent with a request to write remote data, this counter is incremented by one.
- The type is int.
- The unit is packets/second.

### **Component: Windows Process**

Use the Process data set to monitor information about a specific process, such as the amount of time the process runs, its thread count, and how it uses real and virtual memory. Process is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### Process Instance Name

- Instance Name. Valid format is a text string of up to 64 characters.
- The type is string. This is a key dimension.

##### Process Parameter

- A special one for the front end to use as a column header.
- The type is string.

##### Process ID Process

- The unique identifier of this process. ID Process numbers are reused, so they only identify a process for the lifetime of that process.
- The type is int. This is a key dimension.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### Process User

- The user ID associated with the running process.
- The type is string.

##### Process Value

- A special one for the front end to use as a column header.
- The type is string.

##### Process Binary Path

- The fully qualified path to the device binary executable running in the process in UTF-8.
- The type is string.

##### Process System Name

- The managed system name.
- The type is string.

#### Process Server Name

- The managed system name.
- The type is string.

### Metrics

#### Process Working Set Peak

- The maximum working set of a process in bytes since the process started.
- The type is double.
- The unit is bytes.

#### Process IO Data Operations per Sec

- The rate the process is issuing read and write I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/Os.
- The type is int.
- The unit is operations/second.

#### Process IO Other Bytes per Sec

- The rate the process is issuing bytes to I/O operations that do not involve data such as control operations. This counter counts all I/O activity generated by the process to include file, network and device I/Os.
- The type is int.
- The unit is bytes/second.

#### Process Page File Bytes

- The number of bytes of page file space a process uses.
- The type is double.
- The unit is bytes.

#### Process Pool Nonpaged Bytes

- The number of bytes of pool nonpaged memory a process uses.
- The type is int.
- The unit is kilobytes.

#### Process Thread Count

- The number of threads currently active in a process.
- The type is int.
- The unit is threads.

#### Process IO Read Bytes per Sec

- The rate the process is reading bytes from I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/Os.
- The type is int.
- The unit is bytes/second.

#### Process Page File kBytes

- The number of KBs of page file space a process uses.
- The type is double.
- The unit is kilobytes.

#### Process Working Set

- The size of the current working set of a process in bytes.
- The type is double.
- The unit is bytes.

#### Process Page Faults/sec

- The average number of page faults that have occurred for a process per second.
- The type is int.
- The unit is faults/second.

#### Process Working Set kBytes Peak

- The maximum number of KBs of a working set.
- The type is double.
- The unit is kilobytes.

#### Process IO Read Operations per Sec

- The rate the process is issuing read I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/Os.
- The type is int.
- The unit is operations/second.

#### Process Page File kBytes Peak

- The maximum number of KBs of page file space a process has used since starting.
- The type is double.
- The unit is kilobytes.

#### Process Virtual kBytes

- The number of KBs of virtual address space that a process uses. Note: the value 9223372036854775807 indicates Value\_Exceeds\_Maximum and the value -9223372036854775808 indicates Value\_Exceeds\_Minimum.
- The type is double.
- The unit is kilobytes.

#### Process Private kBytes

- The number of KBs of memory space a process has allocated that cannot be shared with other processes.
- The type is double.
- The unit is kilobytes.

#### Process Virtual Bytes

- The number of bytes of virtual address space that a process uses.
- The type is double.
- The unit is bytes.

#### Process Priority Base

- The current base priority of a process.
- The type is int.
- The unit is priority.

#### Process % User Time

- The percentage of elapsed time that a process has executed instructions in user mode. Valid values are positive integers that can include the use of the \*AVG, \*MIN, \*MAX, or \*SUM functions. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### Process Virtual Bytes Peak

- The maximum number of bytes of virtual address space a process has used since starting.
- The type is double.
- The unit is bytes.

#### Process Virtual kBytes Peak

- The maximum number of KBs of virtual address space a process has used since starting.
- The type is double.
- The unit is kilobytes.

#### Process Private Bytes

- The number of bytes of memory space a process has allocated that cannot be shared with other processes.
- The type is double.
- The unit is bytes.

#### Process Pool Paged Bytes

- The number of bytes of pool paged memory a process uses.
- The type is int.
- The unit is bytes.

#### Process % Privileged Time

- The percentage of elapsed time that a process has executed instructions in privileged mode. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

#### Process Page File Bytes Peak

- The maximum number of bytes of page file space a process has used since starting.
- The type is double.
- The unit is bytes.

#### Process IO Write Operations per Sec

- The rate the process is issuing write I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/Os.
- The type is int.
- The unit is operations/second.

#### Process % Processor Time

- The percentage of elapsed time that a process has used the processor to execute instructions.
- The type is int.
- The unit is percent.

#### Process Count

- The count of process executable instances.
- The type is int.
- The unit is instances.

#### Process IO Write Bytes per Sec

- The rate the process is writing bytes to I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/Os.
- The type is int.
- The unit is bytes/second.

#### Process Handle Count

- The total number of handles currently open through this process.
- The type is double.
- The unit is handles.

#### Process IO Data Bytes per Sec

- The rate the process is reading and writing bytes in I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/Os.
- The type is int.
- The unit is bytes/second.

#### Process Working Set kBytes

- The size of the current working set of a process in KBs.
- The type is double.
- The unit is kilobytes.

#### Process IO Other Operations per Sec

- The rate the process is issuing I/O operations that are neither a read or a write operation. An example of this type of operation would be a control function. This counter counts all I/O activity generated by the process to include file, network and device I/Os.



- The type is int.
- The unit is operations/second.

#### Process Elapsed Time (Seconds)

- The total amount of time, in seconds, a process has been running.
- The type is double.
- The unit is seconds.

#### Process Avg % Processor Time

- Percentage of processor use, as an average across all processors in the system. Note that the attribute value is averaged in synch with the situation or historical collection interval.
- The type is int.
- The unit is percent.

### **Component: Active Tasks running on VCenter Server**

Use the Print Queue data set to create situations that monitor the performance and operation of printers locally attached to a computer. Print Queue is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set. Attribute values in this table are obtained from PerfMon. All attribute values are provided for printers locally attached to the computer. Network printers, file printers, and printers attached to remote print servers do not have all their values shown in the local computer's PerfMon database. For these printers, some metrics, such as Job Errors, Out of Paper Errors, Not Ready Errors, are reported as zero.

#### **Dimensions**

##### Print Queue Collection Time

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### Print Queue Name (Unicode)

- Name of the print queue object for monitoring print server activity.
- The type is string. This is a key dimension.

##### Print Queue Name

- Name of the print queue object for monitoring print server activity.
- The type is string.

##### node

- The managed system name. .
- The type is string.

#### **Metrics**

##### Print Queue Bytes Printed/sec (Superseded)

- Number of bytes per second printed on a print queue.
- The type is int.
- The unit is bytes/second.

##### Print Queue Jobs Spooling

- Current number of spooling jobs in a print queue.
- The type is int.

- The unit is jobs.

#### Print Queue Not Ready Errors

- Total number of printer not ready errors in a print queue since the last restart.
- The type is int.
- The unit is errors.

#### Print Queue Out of Paper Errors

- Total number of out of paper errors in a print queue since the last restart.
- The type is int.
- The unit is errors.

#### Print Queue Total Pages Printed

- Total number of pages printed through GDI on a print queue since the last restart.
- The type is double.
- The unit is pages.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### Print Queue Max References

- Peak number of references (open handles) to this printer.
- The type is int.
- The unit is handles.

#### Print Queue Average Out Of Paper Errors/Day

- Average number of out of paper errors per day, where a day is a complete 24 hour period, since the system was last started. Calculated as Out Of Paper Errors System Up Time Days. If the system has been running less than 1 day, this value is the same as Out Of Paper Errors.
- The type is int.
- The unit is errors/day.

#### Print Queue Enumerate Network Printer Calls

- Total number of calls from browse clients to this print server to request network browse lists since last restart.
- The type is int.
- The unit is calls.

#### Print Queue Average Job Errors/Day

- Average number of job errors per day, where a day is a complete 24 hour period, since the system was last started. Calculated as Job Errors System Up Time Days. If the system has been running less than 1 day, this value is the same as Job Errors.
- The type is int.
- The unit is jobs/day.

#### Print Queue Max Jobs Spooling

- Maximum number of spooling jobs in a print queue since last restart.
- The type is int.
- The unit is jobs.

#### Print Queue Add Network Printer Calls

- Total number of calls from other print servers to add shared network printers to this server since last restart.
- The type is int.
- The unit is calls.

#### Print Queue Bytes Printed/sec

- Number of bytes per second printed on a print queue. This attribute is the 64-bit version of Bytes Printed/sec.
- The type is double.
- The unit is bytes/second.

#### Print Queue Average Not Ready Errors/Day

- Average number of not ready errors per day, where a day is a complete 24 hour period, since the system was last started. Calculated as Not Ready Errors System Up Time Days. If the system has been running less than 1 day, this value is the same as Not Ready Errors.
- The type is int.
- The unit is errors/day.

#### Print Queue Job Errors

- Total number of job errors in a print queue since last restart.
- The type is int.
- The unit is errors.

#### Print Queue Total Jobs Printed

- Total number of jobs printed on a print queue since the last restart. This attribute is the 64-bit version of Total Jobs Printed.
- The type is double.
- The unit is jobs.

#### Print Queue References

- Current number of references (open handles) to this printer.
- The type is int.
- The unit is handles.

#### Print Queue Total Pages Printed (Superseded)

- Total number of pages printed through GDI on a print queue since the last restart.
- The type is int.
- The unit is pages.

#### Print Queue Total Jobs Printed (Superseded)

- Total number of jobs printed on a print queue since the last restart. This attribute is the 64-bit version of Total Pages Printed.

- The type is int.
- The unit is jobs.

#### Print Queue Jobs

- Current number of jobs in a print queue.
- The type is int.
- The unit is jobs.

### **Component: Active Tasks running on VCenter Server**

Use the Services data set to obtain status and configuration information about all of the services installed on the Windows Server. Services are background processes run by the operating system, regardless of the user logged in to the system. Services is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

#### **Dimensions**

##### Services Current State

- The current state of the service. This state can be Stopped, Start Pending, Stop Pending, Running, Continue Pending, Paused Pending, Paused, or Unknown. Valid format is a text string of up to 20 characters. For example, Running indicates that the service is currently running.
- The type is string.

##### Service Dependency

- The name of a service or load order group that must start before the given service can start. If there are no dependencies for the given service, this field is blank. For example, +TID indicates the name of a load order group that must start first.
- The type is string. This is a key dimension.

##### Service Name (Unicode)

- The internal name of the service in the Service Control Manager database. The maximum size of the string is 64 bytes.
- The type is string. This is a key dimension.

##### Services Load Order Group

- The name of the load ordering group of which this service is a member. Services can be placed in groups so other services can have dependencies on a group of services. If the service is not in a load ordering group, then this field is blank. For example, Network Provider is an example of a load order group.
- The type is string.

##### Services Display Name (Unicode)

- The name of the service as it is displayed in the NT Service Control Manager applet in UTF8.
- The type is string.

##### Services Start Type

- Specifies how to start the service, including Boot, System, Automatic, Manual, Disabled, Unknown, and Delayed (on Windows 2008 or later systems).
- The type is string.

##### Service Name

- The internal name of the service in the Service Control Manager database. Valid format is a text string of up to 256 characters. For example, NWCWorkstation is an example of a service name.
- The type is string. This is a key dimension.

node

- The managed system name.
- The type is string.

Services Binary Path (Unicode)

- The fully qualified path to the service binary executable in UTF8.
- The type is string.

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

Services Account Id

- The account name under which the service process is logged on when it runs. It takes the form of DomainNameUserName such as LocalSystem.
- The type is string.

Services Account ID (Unicode)

- The account name under which the service process is logged on when it runs in UTF8. It takes the form of DomainNameUserName such as . LocalSystem. Valid format is a text string of up to 52 bytes.
- The type is string.

Services Path to Executable

- The fully qualified path to the service binary executable. For example, D:WINNTSystem32Services.exe indicates the path to the service binary executable.
- The type is string.

Services Display Name

- The name of the service as it is displayed in the Service Control Manager applet. Valid format is a text string of up to 64 characters. For example, Gateway Service for Network is an example of a display name.
- The type is string.

## Metrics

Row Number

- Row Number.
- The type is int.
- The unit is row.

Services Trigger Events

- The number of triggered events of the service.
- The type is int.
- The unit is events.

## Component: Active Tasks running on VCenter Server

Use the Thread data set to monitor information about a specific threads within a process, such as the amount of time the thread runs, its CPU usage, and its state. Thread is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

### Dimensions

#### Thread Collection Time

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### Thread Parameter

- A special one for the front end to use as a column header.
- The type is string.

#### node

- The managed system name.
- The type is string.

#### Thread Base Thread Priority

- The current base priority of this thread. The system can raise and lower a thread's base priority relative to the process base priority.
- The type is int.

#### Thread Wait Reason

- Thread Wait Reason is only applicable when the thread is in the Wait state (see Thread State). It is 0 or 7 when the thread is waiting for the Executive, 1 or 8 for a Free Page, 2 or 9 for a Page In, 3 or 10 for a Pool Allocation, 4 or 11 for an Execution Delay, 5 or 12 for a Suspended condition, 6 or 13 for a User Request, 14 for an Event Pair High, 15 for an Event Pair Low, 16 for an LPC Receive, 17 for an LPC Reply, 18 for Virtual Memory, 19 for a Page Out; 20 and higher are not assigned at the time of this writing.
- The type is int.

#### Thread Process Thread Id

- Instance Name for thread (Process/ThreadID).
- The type is string. This is a key dimension.

#### Thread Virtual Address

- The starting virtual address for this thread.
- The type is int.

#### Thread State

- The current state of the thread. It is 0 for Initialized, 1 for Ready, 2 for Running, 3 for Standby, 4 for Terminated, 5 for Wait, 6 for Transition, 7 for Unknown.
- The type is int.

#### Thread Identifier

- The unique identifier of this thread. ID Thread numbers are reused, so they only identify a thread for the lifetime of that thread.
- The type is int.

#### Thread Value

- A special one for the front end to use as a column header.
- The type is string.

#### Thread Process Identifier

- The unique identifier of this process. ID Process numbers are reused, so they only identify a process for the lifetime of that process.
- The type is int.

#### Thread Current Thread Priority

- The current priority of this thread.
- The type is int.

### **Metrics**

#### Row Number

- Row Number.
- The type is int.
- The unit is row.

#### Thread Elapsed Time Used Percent

- The percentage of elapsed time that this thread used the processor to execute instructions.
- The type is int.
- The unit is percent.

#### Thread Context Switches

- The rate of switches from one thread to another.
- The type is int.
- The unit is switches/second.

#### Thread Privileged Elapsed Time Percent

- The percentage of elapsed time that this thread has spent executing code in privileged mode.
- The type is int.
- The unit is percent.

#### Thread Elapsed Time (Seconds)

- The total elapsed time (in seconds) this thread has been running.
- The type is int.
- The unit is seconds.

#### Thread User Mode Elapsed Time Percent

- The percentage of elapsed time that this thread has spent executing code in user mode.
- The type is int.
- The unit is percent.

## Component: Active Tasks running on VCenter Server

Use the SMTP (Simple Mail Transfer Protocol) Server data set to create situations to monitor a wide range of activities associated with the hosting of an electronic mail server. SMTP Server is a multiple-instance data set. You cannot mix these attributes with those of any other multiple-instance data set.

### Dimensions

SMTP Server

- Instance name of SMTP virtual server.
- The type is string. This is a key dimension.

Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

node

- The managed system name.
- The type is string.

### Metrics

SMTP Server kBytes Total

- The total number of KBs sent and received.
- The type is int.
- The unit is kilobytes.

SMTP Server Messages Delivered Total

- The total number of messages delivered to local mailboxes.
- The type is int.
- The unit is messages.

SMTP Server Total Connection Errors

- The total number of connection errors.
- The type is int.
- The unit is errors.

SMTP Server Message Delivery Retries

- The total number of local deliveries that were retried.
- The type is int.
- The unit is retries.

SMTP Server Message Bytes Sent/sec

- The rate that bytes are sent in messages.
- The type is int.
- The unit is bytes/second.

SMTP Server Messages Received/sec

- The rate that inbound messages are being received.
- The type is int.



- The unit is messages/second.

#### SMTP Server Directory Pickup Queue Length

- The number of messages in the directory pickup queue. Note that this attribute is not available on systems running Windows 2000 with IIS 5.0.
- The type is int.
- The unit is messages.

#### SMTP Server Connection Errors/sec

- The number of connection errors per second.
- The type is int.
- The unit is errors/second.

#### SMTP Server Directory Drops Total

- The total number of messages placed in a drop directory.
- The type is int.
- The unit is messages.

#### SMTP Server Routing Table Lookups Total

- The total number of routing table lookups. Note: the value 2147483647 indicates Value\_Exceeds\_Maximum and the value -2147483648 indicates Value\_Exceeds\_Minimum.
- The type is int.
- The unit is lookups.

#### SMTP Server ETRN Messages Total

- The total number of ETRN messages received by the server.
- The type is int.
- The unit is messages.

#### SMTP Server Outbound Connections Total

- The total number of outbound connections attempted.
- The type is int.
- The unit is connections.

#### SMTP Server Bytes Sent/sec

- The rate that bytes are sent.
- The type is int.
- The unit is bytes/second.

#### SMTP Server Number of MailFiles Open

- Number of handles to open mail files.
- The type is int.
- The unit is handles.

#### SMTP Server Directory Drops/sec

- The number of messages placed in a drop directory per second.
- The type is int.

- The unit is messages/second.

#### SMTP Server Messages Sent Total

- The total number of outbound messages sent.
- The type is int.
- The unit is messages.

#### SMTP Server kBytes Sent Total

- The total number of KBs sent.
- The type is int.
- The unit is kilobytes.

#### SMTP Server NDRs Generated

- The number of non-delivery reports that have been generated.
- The type is int.
- The unit is reports.

#### SMTP Server % Recipients Local

- The percentage of recipients that is delivered locally.
- The type is int.
- The unit is percent.

#### SMTP Server Local Queue Length

- The number of messages in the local queue.
- The type is int.
- The unit is messages.

#### SMTP Server Messages Delivered/sec

- The rate that messages are delivered to local mailboxes.
- The type is int.
- The unit is messages/second.

#### SMTP Server % Recipients Remote

- The percentage of recipients that is delivered remotely.
- The type is int.
- The unit is percent.

#### SMTP Server Message kBytes Sent Total

- The total number of KBs sent in messages.
- The type is int.
- The unit is kilobytes.

#### SMTP Server DNS Queries Total

- The total number of DNS lookups.
- The type is int.
- The unit is lookups.

#### SMTP Server Message Send Retries

- The total number of outbound message sends that were retried.
- The type is int.
- The unit is messages.

#### SMTP Server Message kBytes Total

- The total number of KBs sent and received in messages.
- The type is int.
- The unit is kilobytes.

#### SMTP Server Avg Retries/msg Delivered

- The average number of retries per local delivery.
- The type is int.
- The unit is retries.

#### SMTP Server Local Retry Queue Length

- The number of messages in the local retry queue.
- The type is int.
- The unit is messages.

#### SMTP Server kBytes Received Total

- The total number of KBs received.
- The type is int.
- The unit is kilobytes.

#### SMTP Server Messages Refused for Address Objects

- The total number of messages refused due to no address objects.
- The type is int.
- The unit is messages.

#### SMTP Server Avg Recipients/msg Received

- The average number of recipients per inbound message received.
- The type is int.
- The unit is recipients.

#### SMTP Server ETRN Messages/sec

- The number of ETRN messages per second.
- The type is int.
- The unit is messages/second.

#### SMTP Server Inbound Connections Current

- The total number of connections currently inbound.
- The type is int.
- The unit is connections.

#### SMTP Server Avg Recipients/msg Sent

- The average number of recipients per outbound messages sent.

- The type is int.
- The unit is recipients.

#### SMTP Server Messages Retrieved/sec

- The rate that messages are being retrieved from the mail pick-up directory.
- The type is int.
- The unit is messages/second.

#### SMTP Server Message kBytes Received Total

- The total number of KBs received in messages.
- The type is int.
- The unit is kilobytes.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like any other attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### SMTP Server Messages Sent/sec

- The rate that outbound messages are being sent.
- The type is int.
- The unit is messages/second.

#### SMTP Server Message Bytes Total/sec

- The rate that bytes are sent and received in messages.
- The type is int.
- The unit is bytes/second.

#### SMTP Server Outbound Connections Current

- The number of connections currently outbound.
- The type is int.
- The unit is connections.

#### SMTP Server Messages Refused for Mail Objects

- The total number of messages refused due to no mail objects.
- The type is int.
- The unit is messages.

#### SMTP Server DNS Queries/sec

- The rate of DNS lookups.
- The type is int.
- The unit is queries/second.

#### SMTP Server Message Bytes Received/sec

- The rate that bytes are received in messages.

- The type is int.
- The unit is bytes/second.

#### SMTP Server Bytes Received/sec

- The rate that bytes are received.
- The type is int.
- The unit is bytes/second.

#### SMTP Server Bytes Total/sec

- The rate that bytes are sent and received.
- The type is int.
- The unit is bytes/second.

#### SMTP Server Avg Retries/msg Sent

- The average number of retries per outbound message sent.
- The type is int.
- The unit is retries.

#### SMTP Server Number of QueueFiles Open

- Number of handles to open queue files.
- The type is int.
- The unit is handles.

#### SMTP Server Remote Queue Length

- The number of messages in the remote queue.
- The type is int.
- The unit is messages.

#### SMTP Server Remote Retry Queue Length

- The number of messages in the retry queue for remote delivery.
- The type is int.
- The unit is messages.

#### SMTP Server Routing Table Lookups/sec

- The number of routing table lookups per second. Note: the value 2147483647 indicates Value\_Exceeds\_Maximum and the value -2147483648 indicates Value\_Exceeds\_Minimum.
- The type is int.
- The unit is lookups/second.

#### SMTP Server Outbound Connections Refused

- The number of outbound connection attempts refused by remote sites.
- The type is int.
- The unit is connections.

#### SMTP Server Inbound Connections Total

- The total number of inbound connections received.
- The type is int.

- The unit is connections.

#### SMTP Server Messages Retrieved Total

- The total number of messages retrieved from the mail pick-up directory.
- The type is int.
- The unit is messages.

#### SMTP Server Messages Received Total

- The total number of inbound messages accepted.
- The type is int.
- The unit is messages.

#### SMTP Server Messages Refused for Size

- The total number of messages rejected because they were too big.
- The type is int.
- The unit is messages.

### **Component: Active Tasks running on VCenter Server**

Use the VMWare Processor data set to create situations that monitor processor statistics for this virtual machine. These attributes are supported only using VMWare ESX 5.0 or later, and only if the VM Tools running on the virtual machine matches the version of the ESX server.

#### **Dimensions**

##### VM Processor Host processor speed in MHz

- The host processor speed.
- The type is int.

##### VM Processor Limit in MHz

- The maximum processing power in MHz allowed to the virtual machine. Assigning a CPU Limit ensures that this virtual machine never consumes more than a certain amount of the available processor power. By limiting the amount of processing power consumed, a portion of the processing power becomes available to other virtual machines.
- The type is int.

##### VM Processor Instance Name

- The instance name.
- The type is string. This is a key dimension.

##### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

##### node

- The managed system name.
- The type is string.

##### VM Processor CPU stolen time

- The time, in ms, that the VM was able to run but not scheduled to run.
- The type is double.

## Metrics

### VM Processor Reservation in MHz

- The minimum processing power in MHz available to the virtual machine. Assigning a CPU Reservation ensures that even as other virtual machines on the same host consume shared processing power, there is still a certain minimum amount for this virtual machine.
- The type is int.
- The unit is megahertz.

### VM Processor Shares

- The number of CPU shares allocated to the virtual machine.
- The type is int.
- The unit is shares.

### VM Processor Effective VM speed in Mhz

- The approximate average effective speed of the VM's virtual CPU during the time period between the two samples.
- The type is int.
- The unit is megahertz.

### VM Processor % Processor Time

- The current load of the VM's virtual processor.
- The type is double.
- The unit is percent.

## Component: DNS Activity

Information about DNS (Domain Name Server) server activity and performance.

### Dimensions

#### DNS WINS DNS Value

- A special one for the front end to use as a column header.
- The type is string.

#### Time Stamp

- The date and time the agent collects information as set on the monitored system.
- The type is timestamp.

#### DNS WINS DNS Parameter

- A special one for the front end to use as a column header.
- The type is string.

#### DNS WINS node

- The managed system name.
- The type is string.

## Metrics

### DNS WINS Reverse Response Sent

- The total number of WINS Reverse lookup responses sent by the server.

- The type is int.
- The unit is responses.

#### DNS WINS Lookup Received

- The total number of WINS lookup requests received by the server.
- The type is int.
- The unit is requests.

#### DNS WINS Response Sent

- The total number of WINS lookup responses sent by the server.
- The type is int.
- The unit is responses.

#### DNS WINS Lookup Received/sec

- The average number of WINS lookup requests received by the server in each second.
- The type is int.
- The unit is requests/second.

#### DNS WINS Reverse Lookup Received

- The total number of WINS reverse lookup requests received by the server.
- The type is int.
- The unit is requests.

#### DNS WINS Reverse Lookup Received/sec

- The average number of WINS reverse lookup requests received by the server in each second.
- The type is int.
- The unit is requests/second.

#### DNS WINS Reverse Response Sent/sec

- The average number of WINS Reverse lookup responses sent by the server in each second.
- The type is int.
- The unit is responses/second.

#### Row Number

- Row number. This attribute is not available for use in situations. Otherwise, this attribute is available to use like anyother attribute, for example it is available for reports, queries, and workspaces.
- The type is int.
- The unit is row.

#### DNS WINS Response Sent/sec

- The average number of WINS lookup responses sent by the server in each second.
- The type is int.
- The unit is responses/second.



## Kubernetes resources, metrics, and events

### Kubernetes and OpenShift Resources, Metrics and Events

#### Overview

The Kubernetes data collector provides metric, resource, and event data on the following resource types. The official type for each entity is listed in parentheses. This information can be used in the creation of event and incident policies and metric thresholds.

Kubernetes Cluster (k8sCluster)

Kubernetes Node (k8sNode)

Kubernetes Service (k8sService)

Kubernetes Ingress (k8sIngress)

OpenShift Route (k8sRoute)

Kubernetes Namespace/OpenShift Project (k8sNamespace)

Kubernetes Resource Quota (k8sResourceQuota)

OpenShift Cluster Resource Quota (k8sClusterResourceQuota)

Kubernetes Pod (k8sPod)

Kubernetes Container (k8sContainer)

Kubernetes Controllers

Kubernetes Deployment & OpenShift DeploymentConfig (k8sDeployment)

Kubernetes Stateful Set (k8sStatefulSet)

Kubernetes Daemon Set (k8sDaemonSet)

Kubernetes Cron Job (k8sCronJob)

Kubernetes Job (k8sJob)

Kubernetes Replica Set (k8sReplicaSet)

Kubernetes Replication Controller (k8sReplicationController)

MCM Subscription (CRD)

Kubernetes-sig Application (CRD)

#### Resources/Metrics

The attributes that are collected for your particular resource are dependent both on what is defined in your object spec, and the Kubernetes version you are running. For more information about Kubernetes object attributes, see the [official Kubernetes documentation](#).

#### Kubernetes Cluster

The following attributes and metrics are collected and monitored for a Kubernetes cluster.

Attributes

Name

Metrics

CPU: Allocatable Nanocores, Capacity Nanocores, Usage Core Nanoseconds, Usage Millicores

Deployment Availability Percent

Ephemeral-Storage: Allocatable Bytes, Capacity Bytes

File system: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes

Hugepages-2Mi: Allocatable Bytes, Capacity Bytes

Memory: Allocatable Bytes, Available Bytes, Capacity Bytes, Major Page Faults, Page Faults, Rss Bytes, Usage Bytes, Usage with Cache Bytes

Pods: Allocatable, Capacity, Hosted

Rlimit: Curproc, Maxpid

Runtime Image file system: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes

Stateful Set: Total StatefulSets, Available Statefulsets, Availability Percent

Deployment: Total Deployments, Deployments Available, Availability Percent

### **OpenShift Cluster Resource Quota**

The following and metrics are collected and monitored for an OpenShift Cluster Resource Quota:

Attributes

Cluster Name

Metric Selector

Name

Namespace

Namespaces

Qualified Name

Selector

Metrics

Configmaps: Hard, Used, Percent Used

CPU Limits: Hard, Used, Percent Used

CPU Requests: Hard, Used, Percent Used

Memory Limits: Hard, Used, Percent Used

Memory Requests: Hard, Used, Percent Used

openshiftio/imagestreams:Hard, Used, Percent Used

Persistent Volume Claims:Hard, Used, Percent Used

Pods:Hard, Used, Percent Used

Replication Controllers:Hard, Used, Percent Used

Storage (Request):Hard, Used, Percent Used

Resource Quotas:Hard, Used, Percent Used

Secrets:Hard, Used, Percent Used

Services:Hard, Used, Percent Used

### **Kubernetes Container**

The following attributes and metrics are collected and monitored for a Kubernetes container:

Cluster Name

Container ID

Image

Image ID

Image Pull Policy

Limits

Name  
Namespace  
Pod Name  
Ports  
Ready  
Termination Message Path  
Termination Message Policy  
Metrics  
CPU: Limits Nanocores, Requests Nanocores, Usage Core Nanoseconds  
Logs: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes  
Memory: Limits Bytes, Major Page Faults, Page Faults, Requests Bytes, Rss Bytes, Usage Bytes, Usage with Cache Bytes  
Restart Count  
Rootfs: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes

### **Kubernetes Cron Job**

The following attributes are collected and monitored for a Kubernetes cron job:

Cluster Name  
Concurrency Policy  
Name  
Namespace  
Qualified name  
Schedule  
Suspend

### **Kubernetes Daemon Set**

The following attributes and metrics are collected and monitored for a Kubernetes Daemon Set:

Annotations  
Cluster Name  
Generation  
Metric Selector  
Name  
Namespace  
Qualified Name  
Selector  
Template Generation  
Updated Number Scheduled  
Update Strategy  
Metrics  
Current Number Scheduled  
Desired Number Scheduled

Number Ready  
Number Available  
Number Misscheduled  
Updated Number Scheduled

### **Kubernetes Deployment/Openshift DeploymentConfig**

The following attributes and metrics are collected and monitored for a Kubernetes Deployment/  
Openshift DeploymentConfig

Attributes

Annotations

Available Replicas

Cluster Name

Generation

Labels

Metric Selector

Name

Namespace

Qualified Name

Selector

Metrics

Available Replicas

Replicas

Unavailable Replicas

Updated Replicas

### **Kubernetes Ingress**

The following attributes are collected and monitored for a Kubernetes Ingress

Attributes

Cluster Name

Generation

Hosts

Labels

Load Balancer

Name

Namespace

Paths

Ports

### **Kubernetes Job**

The following attributes and metrics are collected and monitored for a Kubernetes Job

Attributes

Backoff Limit

Cluster Name  
Completion Time  
Labels  
Metric Selector  
Name  
Namespace  
Parallelism  
Qualified Name  
Selector  
Start Time  
Metrics  
Failed  
Succeeded  
Completions

### **Kubernetes Namespace/OpenShift Project**

The following attributes and metrics are collected and monitored for a Kubernetes Namespace/OpenShift Project

Attributes

Cluster Name

Name

Metrics

CPU Usage Core Nanoseconds

Ephemeral Storage Used Bytes

Memory: Available Bytes, Major Page Faults, Page Faults, Usage Bytes, Working Set Bytes,

Network: Received Bytes, Received Errors, Transmitted Bytes, Transmitted Errors

Num Containers

Restart Count

### **Kubernetes Node**

The following attributes and metrics are collected and monitored for a Kubernetes Node

Attributes

Allocatable

Annotations

Architecture

Boot ID

Capacity

Cluster Name

Container Runtime Version

External ID

Hostname

Internal IP

Kernel Version

Kube Proxy Version

Kubelet Port

Kubelet Version

Labels

MachineID

Name

Node Role

Operating System

OS Image

Pod CIDR

System UUID

Unschedulable

Volumes In Use

Metrics

CPU: Allocatable Nanocores, Capacity Nanocores, Usage Core Nanoseconds, Usage Millicores

Ephemeral-Storage: Allocatable Bytes, Capacity Bytes

File System: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes

Hugepages-2Mi: Allocatable Bytes, Capacity Bytes

Memory: Allocatable Bytes, Available Bytes, Capacity Bytes, Major Page Faults, Page Faults, Rss Bytes, Usage Bytes, Usage with Cache Bytes

Pods: Allocatable, Capacity, Hosted

Rlimit: Curproc, Maxpid

Runtime Image File System: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes

### **Kubernetes Pod**

The following attributes and metrics are collected and monitored for a Kubernetes Pod

Attributes

Active Deadline Seconds

Annotations

Automount Service Account Token

Cluster Name

Dns Policy

Generate Name

Host IP

Host IPC

Host Network

Host PID

Hostname  
Image Pull Secrets  
Labels  
Name  
Namespace  
Node Name  
Node Selector  
Phase  
Pod IP  
Priority  
Priority Class Name  
Qos Class  
Readiness Gates  
Restart Policy  
Scheduler Name  
Service Account  
Service Account Name  
Share Process Namespace  
Start Time  
Subdomain  
Termination Grace Period Seconds  
Metrics  
CPU: Usage Core Nanoseconds  
Ephemeral-Storage: Available Bytes, Capacity Bytes, Inodes, Inodes Free, Inodes Used, Used Bytes  
Memory: Major Page Faults, Page Faults, Rss Bytes, Usage Bytes, Usage with Cache Bytes  
Network: Received Bytes, Received Errors, Transmitted Bytes, Transmitted Errors  
Num Containers  
Restart: Count  
Termination Grace Period Seconds

### **Kubernetes Replica Set**

The following attributes and metrics are collected and monitored for a Kubernetes Replica Set

Attributes  
Annotations  
ClusterName  
Generation  
Labels  
Metric Selector  
Name  
Namespace

Qualified Name  
Selector  
Metrics  
Fully Labeled Replicas  
Observed Generation  
Ready Replicas  
Replicas

### **Kubernetes Replication Controller**

The following attributes and metrics are collected and monitored for a Kubernetes Replication Controller

Attributes  
Cluster Name  
Creation Timestamp  
Labels  
Metric Selector  
Name  
Namespace  
Qualified Name  
Selector  
Metrics  
Fully Labeled Replicas  
Observed Generation  
Ready Replicas  
Replicas

### **Kubernetes Resource Quota**

Attributes  
Cluster Name  
Name  
Namespace  
Qualified Name  
Scopes  
Metrics  
Configmaps: Hard, Used, Percent  
CPU Limits: Hard, Used, Percent  
CPU Requests: Hard, Used, Percent  
Memory Limits: Hard, Used, Percent  
Memory Requests: Hard, Used, Percent  
openshiftio/imagestreams: Hard, Used, Percent  
Persistent Volume Claims: Hard, Used, Percent



Pods: Hard, Used, Percent

Replication Controllers: Hard, Used, Percent

Storage (Request): Hard, Used, Percent

Secrets: Hard, Used, Percent

Services: Hard, Used, Percent

### **OpenShift Route**

The following attributes are collected and monitored for an OpenShift Route

Attributes

Annotations

Cluster Name

Creation Timestamp

Host

Labels

Name

Namespace

Qualified Name

Target Port

Wildcard Policy

### **Kubernetes Service**

The following attributes and metrics are collected and monitored for a Kubernetes Service

Attributes

Cluster Name

Cluster IP

External Name

External IPs

External Traffic Policy

Health Check Node Port

Labels

Load Balancer IP

Load Balancer Source Ranges

Merge Tokens

Name

Namespace

Ports

Publish Not Ready Addresses

Selector

Service Type

Session Affinity

Request Path

Request Type  
Metrics  
Browser: Load Time (ms)  
Content Loading Time (ms)  
Error Count per Interval  
Latency (ms)  
Page Transfer Time (ms)  
Real User Latency (ms)  
Resolve Time (ms)  
Status Code

### **Kubernetes Stateful Set**

The following attributes and metrics are collected and monitored for a Kubernetes Stateful Set

Attributes  
Annotations  
Cluster Name  
Labels  
Metric Selector  
Name  
Namespace  
Pod Management Policy  
Qualified Name  
Revision History Limit  
Selector  
Service Name  
Update Revision  
Update Strategy  
Metrics  
Collision Count  
Current Replicas  
Ready Replicas  
Replicas

### **Kubernetes Subscription**

The following attributes are collected and monitored for a Kubernetes Subscription

Attributes  
Annotations  
Channel  
Cluster Name  
Generation  
Labels

Last Update Time  
Name  
Namespace  
Qualified Name  
Source  
Source Namespace

### **Kubernetes Application CRD**

The following attributes are collected and monitored for a Kubernetes Application

Attributes  
Assembly Phase  
Cluster Name  
Component Kinds  
Description  
Generation  
Icons  
Keywords  
Links  
Maintainers  
Name  
Namespace  
Owners  
Qualified Name  
Selector  
Type  
Version

### **Eventing**

#### **Events**

The Kubernetes Monitor **forwards all events raised by Kubernetes** to the Cloud Event Management (CEM) Service. For resource types that are collected by the monitor (as listed above), an ICAM dashboard drilldown URL will be provided in the event payload. The following are the 'action-needed' event types currently raised by Kubernetes (for more information see Kubernetes documentation):

Unhealthy (readiness/liveness failures)  
Backoff  
Conflict  
Error  
Failed  
FailedCreate  
FailedDelete  
FailedMount

FailedSync  
FailedValidation  
FreeDiskSpaceFailed  
HostPortConflict  
HostNetworkNotSupported  
InsufficientFreeCPU  
InsufficientFreeMemory  
InvalidDiskCapacity  
KubeletSetupFailed  
NodeNotReady  
NodeSelectorMismatching  
NodeNotSchedulable  
NodeOutOfDisk  
CIDRNotAvailable  
CIDRAssignmentFailed  
InvalidDiskCapacity  
FreeDiskSpaceFailed  
OutOfDisk

**Note:** Beyond these types, Kubernetes also sends events for 'activities' (or 'normal' events') such as image pulls, container creates, scaled deployments, etc. These 'normal' events will be set with `resolution=False`, so they do not require action by the user, but can be used to provide context to the underlying problem (for example, a rise in application memory usage after updating the Deployment with an image that contains a memory leak)

### Event Grouping/Composition of Incidents

The K8Monitor groups events using the 'qualifiedName' of the highest level pod controller (i.e. Pod/ReplicaSet/Deployment, Pod/StatefulSet, Pod/DaemonSet) under the 'controller' correlation field, if applicable, or by `sourceId`, otherwise. Here, `qualifiedName` takes the form 'EntityName (Namespace:ClusterName)'. This grouping mechanism allows events raised from lower level resources to be grouped under a single incident representing the high level controller, greatly reducing the amount of 'noise' on the incident view.

All events sourced from a particular resource will appear on that resource's timeline. Beyond this, we also propagate a subset of events to their related resources timelines. This is a powerful and unique feature that allows users to see how related components and their health/activities may be impacting a resource. A summary of the current propagation process is as follows:

Cluster timeline: all Node and Pod Controller (Deployment, ReplicaSet, Replication Controller, Job, CronJob, StatefulSet, DaemonSet) events except 'Normal' CronJob events (to prevent noise)

Node timeline: target Node and all Pod Controller (Deployment, ReplicaSet, ReplicationController, Job, CronJob, StatefulSet, DaemonSet) events except 'Normal' CronJob events (to prevent noise)

Service timeline: target Service and all related Pod events

Namespace/Project timeline: events for all resources belonging to the namespace, except 'Normal' Pod events (such as image pull and container create) to prevent noise

Propagating events has the additional effect of propagating the status of the event's source (if it is more severe). For example, say we have a Node 'Critical' event that is propagated to the Cluster timeline—the Cluster entity will adopt this 'Critical' status until the event is cleared.

## Event Severity Transformation

Currently, there are only two event severities defined by the Kubernetes API- "Warning" and "Normal". Since not all 'Warning' events are created equal, **we employ domain level knowledge to further map the severity, thus allowing more serious events to bubble to the top of the incident list.** The current process used to do so is as follows:

Node 'Warning' events: promote to 'Major'

Pod 'Warning' events with 'Fail' in 'reason' (i.e. 'FailedScheduling', 'FailedMount', etc.): promote to 'Minor'

If a user would like to expand on this to further enrich events/incidents, they may do so using an **Event or Incident Policy** (more details in examples below)

## Default Expirations

The K8Monitor populates the 'expiry' field to **10800 seconds (3 hours), as default.**

## Thresholds and Policies

**The K8Monitor does not currently come with any out of the box thresholds or policies. This will likely be added in future releases,** but please note that for many of these resources, it is not required to set manual thresholds for ICAM to detect and create incidents for problem scenarios. For example, if a Deployment is unable to become ready due to an ImagePullBackOff, or it is unable to be scheduled due to insufficient resources, Kubernetes will raise an event that will automatically be forwarded to ICAM. Thresholds (and/or policies) are a good tool for providing finer grained control for (or escalating) alert types that may be of particular importance to your business.

## Support threshold definitions

Kubernetes Cluster (k8sCluster)

Kubernetes Container (k8sContainer)

Kubernetes Daemon Set (k8sDaemonSet)

Kubernetes Deployment (k8sDeployment)

Kubernetes Job (k8sJob)

Kubernetes Replication Controller (k8sReplicationController)

Kubernetes Replica Set (k8sReplicaSet)

Kubernetes Stateful Set (k8sStatefulSet)

Kubernetes Node (k8sNode)

Kubernetes Pod (k8sPod)

Kubernetes Service (k8sService)

NOTE: Some metrics cannot be used in a threshold definition with multiple AND conditions:

Request Name and Latency (ms) with Cluster Ip, Creation Timestamp, or Error Count per Interval

Labels, Latency (ms), Load Balancer, Name, Namespace, Ports, Request Name

## Examples



---

# Chapter 21. Upgrading

Upgrade your Cloud App Management server, agents, and ICAM Data Collectors to get the latest features and functions that are available in the current release.

You can upgrade from the IBM Cloud App Management, Advanced V2019.3.0 offering to the IBM Cloud App Management, Advanced V2019.4.0 offering. For more information, see [Chapter 12, “Upgrading the Cloud App Management server from V2019.3.0 to V2019.4.0,” on page 189.](#)

## Upgrading your agents

Learn how to upgrade your agents and view monitoring data on the Cloud App Management console.

If you have ICAM Agents agents connecting to the Cloud APM server, you can upgrade these agents and view monitoring data on the Cloud App Management console. For more information, see [“Upgrading your ICAM Agents” on page 1387.](#)

If you have IBM Tivoli Monitoring agents or ITCAM agents (referred to as V6 agents) connecting to the Cloud APM server, you can upgrade these agents and view monitoring data on the Cloud App Management console. For more information, see [“Upgrading your IBM Tivoli Monitoring agents” on page 1388.](#)

If you have Cloud APM, Private V8.1.4 agents (referred to as V8 agents) connecting to the Cloud APM server, you can upgrade these agents and view monitoring data on the Cloud App Management console. For more information, see [“Upgrading your Cloud APM agents” on page 1391.](#)

## Upgrading your ICAM Agents

Periodically, new archive files that contain upgraded monitoring agents are available for download. Archive files are available from [IBM Passport Advantage.](#)

### About this task

To upgrade your ICAM Agents, complete the following steps:

### Procedure

1. Download the compressed installation images for the ICAM Agents from the [IBM Passport Advantage](#) website. For more information, see [“Downloading agents and data collectors from Passport Advantage” on page 194.](#)
2. Configure the installation images for communication with the Cloud App Management server. For more information, see [“Configuring the downloaded images” on page 194.](#)
3. Upgrade the agent. For more information, see the following topics:
  - [“Installing agents on UNIX systems” on page 196](#)
  - [“Installing agents on Linux systems” on page 200](#)
  - [“Installing agents on Windows systems” on page 204](#)

**Note:** Agent specific configuration is retained on upgrade.

### Results

The agent is upgraded to the latest version.

### What to do next

Log in to the Cloud App Management console to verify that your agent still reports data.

## Agents on AIX: Stopping the agent and running `slibclean` before you upgrade

If you are upgrading an agent as a non-root user on AIX systems, you must complete this task. Before you run the agent installer, you must stop the agent and run `slibclean` to clear the `libkududp.a` library.

### Procedure

1. Stop the agent by running one of the following commands, depending on whether the agent supports multiple instances:

- ```
./name-agent.sh stop
```

- ```
./name-agent.sh stop instance_name
```

See “Using agent commands” on page 226.

2. Run the following command with root user privileges.

```
slibclean
```

See [slibclean Command](#) in the IBM Knowledge Center.

### Results

The agent is stopped and the `libkududp.a` library is cleared.

### What to do next

Run the agent installer to upgrade the agent to the release that you have downloaded. If the upgrade fails, reboot the server and repeat the procedure.

## Upgrading your IBM Tivoli Monitoring agents

If your IBM Tivoli Monitoring agents or ITCAM agents (referred to as V6 agents) are connected to the Cloud APM server, you can upgrade these agents and view monitoring data on the Cloud App Management console.

### Before you begin

Make sure that you apply the correct version of agent patch. Apply `6.3.0.7-TIV-ITM_TEMA-IF0003` or a later patch if you want to connect the IBM Tivoli Monitoring agents to Cloud App Management server over HTTP. Apply `6.3.0.7-TIV-ITM_TEMA-IF0008` or a later patch if you want to connect the IBM Tivoli Monitoring agents to Cloud App Management server over HTTPS.

### About this task

Complete the following steps to upgrade your V6 agents. If your agent patch is already the correct version, you can skip step 1 and 2, and directly perform step 3.

### Procedure

1. Download the correct agent patch.

- a) Download the agent patch from [IBM Fix Central](#):

- To connect the IBM Tivoli Monitoring agents to Cloud App Management server over HTTP, download `6.3.0.7-TIV-ITM_TEMA-IF0003` or a later patch.
- To connect the IBM Tivoli Monitoring agents to Cloud App Management server over HTTPS, download `6.3.0.7-TIV-ITM_TEMA-IF0008` or a later patch.

- b) **Local configuration only:** Determine the architecture of the target operating system to select the appropriate patch file to apply.

**Tip:** Use the `install_dir/bin/cinfo` script to get the architecture code of the operating system.

- c) **Remote configuration only:** Make sure the OS agent is installed on the remote system.



Be aware of the following limitations before you proceed to apply the agent patch.

**Known limitation:**

- (WebSphere Applications agent) Transaction tracking data is not yet supported by Cloud App Management. If the V6 agent has been enabled for transaction tracking data collection, reconfigure the V6 agent to disable it before you connect the V6 agent to the Cloud App Management server. For more information, see the V6 agent documentation.
  - The agent patch cannot be applied on the system where the monitoring server or portal server is also installed.
  - After the agent patch is applied, the agent subscription facility (ASF) is started. Many ASF related activities might be logged. You can ignore these messages in logs and no action is required.
2. Follow instructions for locally applying the agent patch or remotely applying the agent patch.
- To locally apply the agent patch, do the steps:

- a. Extract the agent patch to the local system where the V6 agent is installed.

In the extracted agent patch directory, different fix files are included for all supported operating systems. Use the appropriate file for the target operating system in the following steps.

- b. Run the following script to apply the patch.

– Linux | UNIX

```
cd temp_dir/agent_patch
./install.sh -h install_dir -q -p `pwd`/unix/tfarch.txt
```

– Windows

```
cd temp_dir\agent_patch\WINDOWS
setup.exe /w /z"/sf%cd%\deploy\TF_Silent_Install.txt" /s
/f2"install_dir\INSTALLITM\Silent_KTF.log"
```


where:

- *temp\_dir* is the temporary directory that contains the extracted agent patch folder.
- *agent\_patch* is the agent patch file name, for example, it is 6.3.0.7-TIV-ITM\_TEMA-IF0003 for connection over HTTP, and 6.3.0.7-TIV-ITM\_TEMA-IF0008 for connection over HTTPS.
- *install\_dir* is the V6 agent installation directory. For example, /opt/ibm/itm.
- *arch* is the architecture code of the operating system. Use the appropriate *tfarch.txt* file for the target system, for example, *tflx8266.txt*.

**Troubleshooting on Windows:** If some product files are locked by other processes on a Windows system, the deployment might fail and the locked files are reported in the `Abort IBM Tivoli Monitoring.log` file.

To solve this problem, manually stop all processes that are locking the files and try again. For example, if you have WebSphere Applications agent installed, you also need to stop the application server that has the agent data collector installed.

Alternatively, you can add `Locked Files=continue` to the installation section in the `TF_Silent_Install.txt` and `TFX64_Silent_Install.txt` files within in the `agent_patch/WINDOWS/Deploy` directory and try again.

For more information about this limitation, see the [Locked files encountered during Windows agent silent installation](#)  technote.

- To remotely apply the agent patch, complete the following steps on a system where the **tacmd** library is available:
  - a. Extract the agent patch to a temporary directory.

There are different `.tar` files for different operating systems in the extracted agent patch directory. Use the appropriate file for the target operating system in the following steps.

- b. On the hub monitoring server system, log in to Tivoli Enterprise Monitoring Server by running the following command from the **tacmd** library:

```
tacmd login -s tems_address -u user_name -p password
```

where:

- *tems\_address* is the host name or IP address of the Tivoli Enterprise Monitoring Server.
- *user\_name* is the user ID that is used to log in to the monitoring server.
- *password* is the user password.

- c. Go to the extracted directory that contains the agent patch for the current operating system.

– **Linux** | **UNIX**

```
cd temp/agent_patch/unix
```

– **Windows**

```
cd temp\agent_patch\WINDOWS/Deploy
```

where:

- *temp* is the temporary directory that contains the extracted agent patch folder.
- *agent\_patch* is the agent patch file name, for example, it is `6.3.0.7-TIV-ITM_TEMA-IF0003` for connection over HTTP, and `6.3.0.7-TIV-ITM_TEMA-IF0008` for connection over HTTPS.

- d. Run the following command to populate the agent depot:

```
tacmd addbundles -i . -t tf
```

After the command is run, more information about the `tf` component, including its version, is returned.

- e. Run the following command from the *tems\_install\_dir/bin* directory to update the agent framework to the version that is returned in step d:

- For connection over HTTP:

```
tacmd updateFramework -n node_name -v 063007003
```

- For connection over HTTPS:

```
tacmd updateFramework -n node_name -v 063007008
```

where, *node\_name* is the node name of the operating system where the V6 agent is installed.

The following example updates the agent framework on the `kvm-011235:LZ` system:

- For connection over HTTP:

```
tacmd updateFramework -n kvm-011235:LZ -v 063007003
```


- For connection over HTTPS:

```
tacmd updateFramework -n kvm-011235:LZ -v 063007008
```

**Troubleshooting on Windows:** If some product files are locked by other processes on a Windows system, the deployment might fail and the locked files are reported in the `Abort IBM Tivoli Monitoring.log` file.

To solve this problem, manually stop all processes that are locking the files and try again. For example, if you have WebSphere Applications agent installed, you also need to stop the application server that has the agent data collector installed.

Alternatively, you can add `Locked Files=continue` to the installation section in the `TF_Silent_Install.txt` and `TFX64_Silent_Install.txt` files within in the `agent_patch/WINDOWS/Deploy` directory and try again.

For more information about this limitation, see the [Locked files encountered during Windows agent silent installation](#)  technote.

3. Follow instructions in [IBM Tivoli Monitoring Knowledge Center](#) to upgrade your V6 agents.
4. Log in to the Cloud App Management console to verify that your agent still reports data. If not, refer to [“Connecting IBM Tivoli Monitoring agents to Cloud App Management server” on page 676](#) to make sure the V6 agents are connected to Cloud App Management server.


## Upgrading your Cloud APM agents

To upgrade Cloud APM agents, you must first connect the agents to Cloud App Management server.

### Before you begin

- Make sure the Cloud APM agents are connected to Cloud App Management server. If not, see [“Connecting Cloud APM agents to Cloud App Management server” on page 690](#) to complete the connecting steps.
- Diagnostics and transaction tracking data are not yet supported by Cloud App Management. If the V8 agents have been enabled for diagnostics and/or transaction tracking data collection, reconfigure the V8 agents to disable them before you connect the V8 agents to the Cloud App Management server.
- For the HTTP Server agent, you must stop the HTTP server before you upgrade the agent.

### Procedure

1. Download the compressed installation images for the ICAM Agents from the [IBM Passport Advantage](#)  website. For more information, see [“Downloading agents and data collectors from Passport Advantage” on page 194](#).
2. Configure the installation images for communication with the Cloud App Management server. For more information, see [“Configuring the downloaded images” on page 194](#).
3. Re-install the agent. For more information, see the following topics:
  - [“Installing agents on UNIX systems” on page 196](#)
  - [“Installing agents on Linux systems” on page 200](#)
  - [“Installing agents on Windows systems” on page 204](#)

### Results

All V8 agents installed on the same system are upgraded. However, you can view monitoring data only for the supported agents on the Cloud App Management console.

### What to do next

Log in to the Cloud App Management console to verify that your agent still reports data.

## Upgrading your data collectors

Periodically, new archive files that contain upgraded ICAM Data Collectors are available for download from [IBM Passport Advantage](#).

### Before you begin

If you have many data collectors installed, you can stagger the updates, for example, to upgrade the data collectors in the Southern region this weekend and the Northern region next weekend. For details, see the Kubernetes tutorial, [Performing a rolling update](#).

Make sure that you upgrade the Cloud App Management server before you upgrade the data collectors. For more information, see [Chapter 12, “Upgrading the Cloud App Management server from V2019.3.0 to V2019.4.0,” on page 189](#).

### Procedure

Complete these steps to upgrade your ICAM Data Collectors:

1. Uninstall your ICAM Data Collectors:
  - To uninstall the Kubernetes data collector, see [“Uninstalling the Kubernetes data collector” on page 574](#).
  - To uninstall the Node.js data collector, see [“Uninstalling the Node.js data collector” on page 605](#).
  - To uninstall the Liberty data collector, see [“Uninstalling the Liberty data collector from your application” on page 599](#).
  - To uninstall the J2SE data collector, see [“Uninstalling the J2SE data collector from your application” on page 589](#).
  - To uninstall the Python data collector, see [“Uninstalling the Python data collector” on page 616](#).
2. Download the installation image and configuration package, configure your ICAM Data Collectors, and validate your re-installation.
  - For instructions about the cloud data collector, see the configuration topics under [“Kubernetes data collector” on page 557](#).
  - For instructions about Node.js data collector, see [“Configuring Node.js application monitoring” on page 601](#).
  - For instructions about Liberty data collector, see [“Configuring Liberty application monitoring” on page 590](#).
  - For instructions about J2SE data collector, see [“Configuring J2SE application monitoring” on page 584](#).
  - For instructions about Python data collector, see [“Configuring Python application monitoring” on page 606](#).

### Results

The ICAM Data Collectors are upgraded to the latest version.

### Special Notice for J2SE data collector

After you upgrade J2SE data collector from 2019.2.0 to a higher version, two J2SE resources with the same name **J2SE Application Runtime** are displayed in the Resources dashboard, one is for 2019.2.0 and the other one is for the upgraded version. Ignore the 2019.2.0 resource and always use the new one to view monitoring statistics. To tell which one is the correct version to view, complete the following steps:

1. Click the **Resources** tab in the Cloud App Management console.
2. Find the J2SE Application Runtime resources on the **Resource groups** page. For more information, see [“Viewing your managed resources” on page 769](#).
3. Check the **Related resources** widget for the J2SE Application Runtime resources. For 2019.2.0, there is no display of **JVM** type. For the updated version, you can see **JVM**.

## Upgrading the Synthetics PoP server

Use the following procedure to upgrade the Synthetics PoP server.

### Procedure

1. Download and unpack the data collectors installation eImage from Passport Advantage (appMgtDataCollectors\_2019.4.0.2.tar.gz). You will see the Synthetics PoP app\_mgmt\_syntheticpop\_xlinux.tar.gz installation file. See [“Downloading agents and data collectors from Passport Advantage”](#) on page 194.
2. Stop the Synthetics PoP and backup the existing Synthetics PoP installation folder.
3. Copy all the files and folders from app\_mgmt\_syntheticpop\_xlinux.tar.gz into the Synthetics PoP installation folder. Over write all files and folders except for:

```
global.environment  
pop.properties  
keyfiles
```

4. Start the Synthetics PoP.



---

## Chapter 22. Troubleshooting and support

Troubleshooting and support information include instructions for resolving problems related to the Cloud App Management product. To resolve Cloud App Management problems, use the following topics to find out the cause of the problem, the symptoms, and how to resolve them. Learn also how to contact IBM support to resolve issues.

For general troubleshooting issues, visit the [dWAnswers](#) forum and for agent-related issues, visit the [Cloud App Management](#) forum.

### Troubleshooting installation and upgrade

Troubleshoot IBM Cloud App Management installation and upgrade issues.

Learn how to isolate and resolve problems with Cloud App Management. Verify that your issues are not related to operating system requirements, such as disk, memory, and CPU capacities. For more information about system requirements for Cloud App Management, see “[System requirements](#)” on page 75. To get support, see [Support](#) in the product documentation.

### Cloud App Management installation fails with permissions error on a Red Hat OpenShift Kubernetes Server (3.11)

The deployment starts then a permissions issue is seen for all pods that use persistent storage.

#### Problem

Installing Cloud App Management on Red Hat OpenShift Kubernetes Server fails with no permissions error on all pods that use persistent storage. Example message: ERROR [main] 2020-02-19 14:49:35,529 CassandraDaemon.java:749 - Has no permission to create directory /opt/ibm/cassandra/bin/./data/data

#### Symptoms

The errors are caused by mounted file system ownership issues.

#### Solution

The stateful sets that use persistent storage need to be modified as follows:

#### Zookeeper

1. Edit the stateful set by running the command:

```
oc edit statefulset ibm-cloud-appmgmt-prod-zookeeper
```

2. Add the following lines into the `initContainers` section:

```
- args:
  - chown 1001:1001 /var/lib/zookeeper/data
  command:
  - /bin/sh
  - -c
  image: alpine:latest
  imagePullPolicy: Always
  name: initcontainer
  resources: {}
  securityContext:
    allowPrivilegeEscalation: false
  capabilities:
    add:
    - CHOWN
    - FOWNER
    - DAC_OVERRIDE
```

```

    drop:
    - ALL
    privileged: false
    readOnlyRootFilesystem: false
    runAsNonRoot: false
    runAsUser: 0
    selinuxOptions:
      type: spc_t
    terminationMessagePath: /dev/termination-log
    terminationMessagePolicy: File
    volumeMounts:
    - mountPath: /var/lib/zookeeper/data
      name: data

```

## Cassandra

1. Edit the stateful set by running the command:

```
oc edit statefulset ibm-cloud-appmgmt-prod-cassandra
```

2. Add the following lines after the dnsPolicy line:

```

initContainers:
  - args:
    - chown 1001:1001 /opt/ibm/cassandra/data && chown 1001:1001 /opt/ibm/cassandra/logs
    command:
    - /bin/sh
    - -c
    image: alpine:latest
    imagePullPolicy: Always
    name: initcontainer
    resources: {}
    securityContext:
      allowPrivilegeEscalation: false
      capabilities:
        add:
        - CHOWN
        - FOWNER
        - DAC_OVERRIDE
      drop:
      - ALL
    privileged: false
    readOnlyRootFilesystem: false
    runAsNonRoot: false
    runAsUser: 0
    selinuxOptions:
      type: spc_t
    terminationMessagePath: /dev/termination-log
    terminationMessagePolicy: File
    volumeMounts:
    - mountPath: /opt/ibm/cassandra/data
      name: data
    - mountPath: /opt/ibm/cassandra/logs
      name: ibm-cloud-appmgmt-prod-cassandralsogs

```

## Couchdb

1. Edit the stateful set by running the command:

```
oc edit statefulset ibm-cloud-appmgmt-prod-couchdb
```

2. Add the following lines after the dnsPolicy line:

```

initContainers:
  - args:
    - chown 1001:1001 /opt/couchdb/data
    command:
    - /bin/sh
    - -c
    image: alpine:latest
    imagePullPolicy: Always
    name: initcontainer
    resources: {}
    securityContext:
      allowPrivilegeEscalation: false

```



```

capabilities:
  add:
    - CHOWN
    - FOWNER
    - DAC_OVERRIDE
  drop:
    - ALL
privileged: false
readOnlyRootFilesystem: false
runAsNonRoot: false
runAsUser: 0
seLinuxOptions:
  type: spc_t
terminationMessagePath: /dev/termination-log
terminationMessagePolicy: File
volumeMounts:
- mountPath: /opt/couchdb/data
  name: data

```

## Kafka

1. Edit the stateful set by running the command:

```
oc edit statefulsets ibm-cloud-appmgmt-prod-kafka
```

2. Add the following lines into the `initContainers` section:

```

- args:
  - chown 1001:1001 /var/lib/kafka/data
  command:
  - /bin/sh
  - -c
  image: alpine:latest
  imagePullPolicy: Always
  name: initcontainer
  resources: {}
  securityContext:
    allowPrivilegeEscalation: false
    capabilities:
      add:
        - CHOWN
        - FOWNER
        - DAC_OVERRIDE
      drop:
        - ALL
    privileged: false
    readOnlyRootFilesystem: false
    runAsNonRoot: false
    runAsUser: 0
    seLinuxOptions:
      type: spc_t
  terminationMessagePath: /dev/termination-log
  terminationMessagePolicy: File
  volumeMounts:
  - mountPath: /var/lib/kafka/data
    name: data

```

## Elasticsearch

1. Edit the stateful set by running the command:

```
oc edit statefulsets ibm-cloud-appmgmt-prod-elasticsearch
```

2. Add the following lines after the `dnsPolicy` line:

```

initContainers:
- args:
  - chown 1000:1000 /opt/elasticsearch/data && sysctl -w vm.max_map_count=262144
  && sysctl -p
  command:
  - /bin/sh
  - -c
  image: alpine:latest
  imagePullPolicy: Always
  name: initcontainer

```

```

resources: {}
securityContext:
  allowPrivilegeEscalation: true
  capabilities:
    add:
      - CHOWN
      - FOWNER
      - DAC_OVERRIDE
    drop:
      - ALL
  privileged: true
  readOnlyRootFilesystem: false
  runAsNonRoot: false
  runAsUser: 0
  seLinuxOptions:
    type: spc_t
terminationMessagePath: /dev/termination-log
terminationMessagePolicy: File
volumeMounts:
- mountPath: /opt/elasticsearch/data
  name: data

```

## Cloud App Management PPA installation fails to load into the IBM Cloud Private repository with Docker authentication error

After you log in to Docker and you start loading the Cloud App Management installation image into the IBM Cloud Private repository, some time later, the loading stops and you get the `unauthorized: authentication required` error.

### Problem

If it takes a substantial amount of time to load the Cloud App Management installation image into the IBM Cloud Private repository, the Docker login times out and the `unauthorized: authentication required` error is displayed. The Cloud App Management installation image stops loading.

### Symptoms

Your images stop loading into IBM Cloud Private repository. Preparing and Waiting messages display in your command window, and then the following error is displayed

```

unauthorized: authentication required
(Are you logged in to the docker registry?)

```

### Solutions

Complete Solution 1 first. If Solution 1 succeeds, you do not need to complete Solution 2. If it fails, you must complete Solution 2.

#### Solution 1

1. Log in to Docker again by issuing the following command:

```
docker login my_cluster_CA_domain:8500
```

where `my_cluster_CA_domain` is the certificate authority (CA) domain, such as `mycluster.icp`. If you did not specify a `my_cluster_CA_domain`, the default value is `mycluster.icp`.

2. If you can successfully log in again, load the Cloud App Management installation image into the IBM Cloud Private repository again by issuing the following command:

```
cloudctl catalog load-archive --archive ppa_file
[--registry my_cluster_CA_domain:8500] [--repo my_helm_repo_name]
```

Where:

- `ppa_file` is the name of the Cloud App Management installation image file.

- *my\_helm\_repo\_name* is the name of the target Helm repository. Run the **cloudctl catalog repos** command to get a list of repositories.

## Solution 2

If the previous Solution 1 fails again and the Cloud App Management PPA installation image stops loading because you are logged out of Docker, complete the following steps:

1. Log in into to Docker again by issuing the following command:

```
docker login my_cluster_CA_domain:8500
```

2. Extract the Cloud App Management installation image by issuing the following command:

```
tar -xvf ppa_file
tar -xvf charts/decompressed_ppa_file
```

where *ppa\_file* is the compressed Cloud App Management PPA installation image file, such as the *icam\_ppa\_2019.4.0\_prod.tar.gz* file.

3. Load all the Docker image archive files into the Docker engine. Record the Docker image name and tag into a file by running the following commands:

```
for i in `ls images/*`;
do
    echo $i;
    echo docker load -i $i
    docker load -i $i | tee -a IMG_LOG.txt
done
```

4. If you are logged out of Docker again, you must log in to Docker again now. Then, tag the Docker images to this tag: *my\_cluster\_CA\_domain:8500/my\_namespace/imagename:tag* by running the following commands:

```
for img in `awk -F ':' ' {print $2:""$3} ' IMG_LOG.txt`
do
    echo docker tag $img my_cluster_CA_domain:8500/my_namespace/$img
    docker tag $img my_cluster_CA_domain:8500/my_namespace/$img
    echo docker push my_cluster_CA_domain:8500/my_namespace/$img
    docker push my_cluster_CA_domain:8500/my_namespace/$img
done
```

where *my\_namespace* is the namespace that the installation image file is loaded to.

5. If you have left your system for more than 12 hours, then you must log in to your IBM Cloud Private cluster again by issuing the following command:

```
cloudctl login -a https://my_cluster_CA_domain:8443 --skip-ssl-validation
```

6. Load the Helm chart that is located in the *charts* directory into IBM Cloud Private by issuing the following command:

```
cloudctl catalog load-chart --archive charts/decompressed_ppa_file
```

## Results

The Cloud App Management PPA installation image is successfully loaded into the IBM Cloud Private repository. You can continue with your Cloud App Management server installation.

## Cloud App Management installation fails with InvalidImageName error

If you use the Helm command line interface (CLI) client to manually install the Cloud App Management server by running the **helm install** command, the installation can fail with the `InvalidImageName` error.

### Problem

The Cloud App Management server installation fails with an `InvalidImageName` error because IBM Cloud Private did not load the Cloud App Management Passport Advantage Archive (PPA) installation image file correctly. This issue occurs because IBM Cloud Private changes the value of the **global.image.repository** setting in the `values.yaml` file.

### Symptoms

When you issue the `kubectl get pod` command, the following output is displayed.

```
kubectl get pod
NAME                                READY   STATUS              RESTARTS   AGE
am-server01-redis-server-554979449d-2n196   0/1     InvalidImageName    0           18m
```

The following output shows some extra trailing characters after the `mycluster.icp:8500/default` value that is incorrect.

```
kubectl describe pod am-server01-redis-server-554979449d-291fk
Name:          am-server01-redis-server-554979449d-291fk
Namespace:    default

Normal       SandboxChanged          18m                kubelet, 9.42.26.89  Pod sandbox changed, it will be killed and re-created.
Warning      Failed                  13m (x26 over 18m) kubelet, 9.42.26.89  Error: InvalidImageName
Warning      InspectFailed           3m (x71 over 18m)  kubelet, 9.42.26.89  Failed to apply default image tag "mycluster.icp:8500/default//redis-ha:4.0.6-r0": couldn't parse image reference "mycluster.icp:8500/default//redis-ha:4.0.6-r0": invalid reference format
```

### Solution

To solve the problem, set the **global.image.repository** setting to the correct value. By completing the following step, you are overriding the value for **global.image.repository** in the `values.yaml` file.

- Issue the following command:

```
helm install mycluster/my_helm_chart_name --set
global.image.repository=my_Cluster_CA_Domain:8500/my_namespace --tls
```

Where

- `my_helm_chart_name` is the name of your Cloud App Management Helm chart, which is `ibm-cloud-appmgmt-prod`.
- `my_Cluster_CA_Domain` is `mycluster.icp` by default.
- `my_namespace` is the namespace that the PPA installation image file is loaded to.

### Results

The Cloud App Management server is successfully installed. When you issue the `kubectl get pod` command to see the status of your pods, the pods display a stable state. Continue with your Cloud App Management deployment.

## Common deployment errors

This topic lists some common errors that you might encounter when you are deploying the Cloud App Management product. The error messages that are displayed, the reasons for these errors occurring, and the solutions for them are included in the following information.

### Problem - Certificate error

Error message such as:

```
could not read x509 key pair (cert: "/root/.helm/cert.pem", key: "/root/.helm/key.pem"):
can't load key pair from cert /root/.helm/cert.pem
and key /root/.helm/key.pem: open /root/.helm/cert.pem: no such file or directory
```

### Symptoms

This error occurs when you try to run any Helm command to connect the Helm server. The certificate key files are not installed correctly.

### Solution

Find the pem keyfiles and copy them to the `/root/.helm` default location. You can manually copy over these keyfiles to the `~/helm/` default directory by issuing the following command:

```
cp ~/.kube/mycluster/*.pem ~/.helm/
```

### Problem - Helm cannot connect error

The error message that is displayed depends on the Helm command that you are running.

```
Error: cannot connect to Tiller
```

```
Error: transport is closing
```

### Symptoms

This error occurs when you try to run a **helm** command and you did not add the `--tls` parameter with this command.

### Solution

Add `--tls` to the **helm** command that you are running. For example:

```
helm install --tls
```

### Problem - kubectl or Helm errors because authorization has expired

For example, error messages such as the following can be displayed

```
kubectl get pods
Error: the server doesn't have a resource type "pods"
```

### Symptoms

These type of errors occur when you are not actively working in your IBM Cloud Private cluster from the IBM Cloud Private CLI for a long time such as a few hours. Your authorization has expired.

## Solution

You must log in to the IBM Cloud Private cluster again by issuing the following command:

```
cloudctl login -a https://mycluster.icp:8443 --skip-ssl-validation -u admin
```

Where `mycluster.icp` is the default IBM Cloud Private cluster name. This name might be something different if you entered your own cluster name when you were deploying IBM Cloud Private.

## Problem - Deployment already exists error

Error message such as:

```
deployment my_deployment_name already exist
```

where `my_deployment_name` is the deployment name.

## Symptoms

This error occurs when the previous deployment with the same name as the deployment you are attempting to install now didn't delete successfully.

## Solution

You must find this deployment and delete it manually. For example, to find the deployment, issue the following command:

```
kubectl get deployment
```

Next, you must manually delete the deployment by issuing the following command:

```
kubectl delete deploy my_deployment_name
```

## 502 Bad Gateway error on IBM Cloud Private

### Problem

A 502 Bad Gateway error is displayed after login to IBM Cloud Private, even though all pods and services are reporting ready. This is an environment configuration issue. The worker nodes have private IP addresses and are unable to communicate with the public IP addresses of the master/proxy node.

### Solution

The worker nodes need access to the public IP addresses of the master/proxy, or the static route must be added.

The following procedure must be performed on each worker node (figures used are for example purposes only):

1. `ip route add 9.46.67.210/32 via 10.21.5.227;ip route add 9.46.67.221/32 via 10.21.5.228`
2. Edit the `/etc/rc.local` file and add the following routes so that they are persisted on restart:

```
ip route add 9.46.67.210/32 via 10.21.5.227
ip route add 9.46.67.221/32 via 10.21.5.228
```

The `9.46.67.210` is the public ip of master, `32` is the network segment. `10.21.5.227` is the private ip of the master.

The 9.46.67.221 is the public ip of proxy, 32 is the network segment. 10.21.5.228 is the private ip of the proxy.

## IBM Multicloud Manager hub cluster fails to import a cluster with Cloud App Management server installed

### Problem

A Red Hat OpenShift cluster with the Cloud App Management server installed prevents other workloads from scheduling.

### Cause

The Cloud App Management server creates a custom SecurityContextConstraint (SCC) when installed in a Red Hat OpenShift cluster with the `security.openshift.io/v1` APIVersion. The created SCC has a pre-defined `priority` value, and may be selected for pods over existing SCCs. For more information, see: [https://access.redhat.com/documentation/en-us/openshift\\_container\\_platform/3.11/html-single/architecture/index#security-context-constraints](https://access.redhat.com/documentation/en-us/openshift_container_platform/3.11/html-single/architecture/index#security-context-constraints)

### Environment

A sample environment where the issue might occur:

Clusters on x86_64	Red Hat OpenShift	IBM Cloud Private	IBM Cloud App Management
IBM Multicloud Manager hub cluster	Red Hat OpenShift V3.11.X	IBM Cloud Private V3.2.1	Cloud App Management server
Remote cluster that is being imported	Red Hat OpenShift V3.11.X	IBM Cloud Private V3.2.1	Cloud App Management server

### Symptom

After the import, the pods on the managed cluster are not running. Failing pods might have the following annotation:

```
openshift.io/scc: ibmcloudappgmt-ibm-cem-ibm-restricted-scc
```

Where `ibmcloudappgmt` is the release name of the Cloud App Management server installation.

### Solution

Delete the custom SecurityContextConstraint (SCC) created during the Cloud App Management server installation. Run the following command:

```
kubectl delete SecurityContextConstraints --selector origin="helm-cem"
```

## When scaling up ibm-redis, the ibm-redis-server pods crash

### Problem

When scaling up `ibm-redis`, the `ibm-redis-server` pods crash and do not become ready.

## Symptom

The following sample environment has 3 sentinels and 2 servers. The master is on `scao-ibm-redis-server-0`, with failover on `scao-ibm-redis-server-1`. After `scao-ibm-redis-server-0` is restarted, both servers show as slaves with no master:

```
freds-mbp:ibm-cem ffabec@us.ibm.com$ kc get po -l redis-role=slave
NAME                                READY    STATUS    RESTARTS   AGE
scao-ibm-redis-server-0             0/1     Running  0          72s
scao-ibm-redis-server-1             1/1     Running  0          2m53s
freds-mbp:ibm-cem ffabec@us.ibm.com$ kc get po -l redis-role=master
No resources found.
freds-mbp:ibm-cem ffabec@us.ibm.com$
```

## Cause

The `ibm-redis-server` pods crash is due to a timing issue in selecting a master.

## Solution

Scale down and up the `ibm-redis-sentinel` or `ibm-redis-server` (or both) StatefulSet to restart it.

```
kubectl scale statefulset 2019.2.1-ibm-redis-server 2019.2.1-ibm-redis-sentinel --replicas=0
kubectl scale statefulset 2019.2.1-ibm-redis-sentinel --replicas=<number> >= 3>
kubectl scale statefulset 2019.2.1-ibm-redis-server --replicas=<number>
```

where:

`2019.2.1` is the release number of the `ibm-redis-server`

`3` is the number of replicas to scale up

## Results

After the StatefulSet is scaled down and then scaled up, it restarts successfully.

## Freeing up resources for large Pod scheduling

Large Pods such as Cassandra fail to deploy due to resources constraints.

### Problem

Large Pods can fail to deploy because resources can run out on a particular node even though the collective Cloud App Management environment together (all nodes added) has these resources available. For example, for the following deployment, a total of 19 cores (5 + 3 + 5 + 6) and 42 GB (9 GB + 13 GB + 10 GB + 10 GB) is available. If you try to deploy a large Pod such as Cassandra that requires four cores and 12 GB, it fails because no single node has these resources available.

```
VM-A: 8 cores 16GB total, 3 core 7GB used, 5 core 9GB available
VM-B: 8 cores 16GB total, 5 core 3GB used, 3 core 13GB available
VM-C: 8 cores 16GB total, 3 core 6GB used, 5 core 10GB available
VM-D: 8 cores 16GB total, 2 core 6GB used, 6 core 10GB available
```

### Symptoms

When you are installing the Cloud App Management server, the `post-install-setup.sh` script displays the list of Pods that are not ready when it is finished running. For more information, see [“Creating your service instance”](#) on page 174.

If this issue occurs, it is probably going to happen for the Cassandra Pod. For example, the script displays the following output for the Cassandra Pod if a node is not available. **IP** is set to `<none>` and **NODE** is set to `<none>`, which indicates that a node that satisfies the requirements cannot be found for Cassandra to



be assigned to. You can also run the **kubectl get pods -o wide** command to display the following output:

POD	READY	STATUS	RESTARTS	AGE	IP	NODE
ibmcloudappgmt-cassandra-0	0/1	ContainerCreating	0	10m	<none>	<none>

## Solution

To deploy Cassandra, you must remove enough resources on a single node to schedule 4 cores and 12 GB.

If you are using local storage with affinity, the resources must be removed from the VM that the PV is linked to. If you are using vSphere, which is storage that is movable, you can select a node that causes the least disruption. For example, select a node that doesn't have any StatefulSet services on it already, or select other nodes with no resources assigned.

Issue the following command to find out which nodes do not have any resources on them:

```
kubectl get pods -o wide --sort-by="{.spec.nodeName}"
```

After a fresh installation of Cloud App Management server, complete the following steps to free up resources on a particular node that you want to use to schedule a large pod such as Cassandra.

1. Scale down all deployments using the following command, this will allow the statefulsets to access space:

```
kubectl scale deploy -l release=<release-name> -n <namespace> --replicas=0
```

2. Once the statefulsets have been assigned, scale back up using the following command:

```
kubectl scale deploy -l release=<release-name> -n <namespace> --replicas=1
```

After an upgrade of Cloud App Management server, complete the following steps to free up resources on a particular node that you want to use to schedule a large pod such as Cassandra.

1. Mark the node so it cannot be scheduled by issuing the following command:

```
kubectl taint node node NoSchedule=true:NoSchedule
```

2. Complete the following steps to delete the Cloud App Management pods from the node until it has enough space for the large pod that you want to schedule. Wait for the deleted pods to be placed on another node.

3. Mark the node so it is available for scheduling again by issuing the following command:

- a. Identify the Pods that you want to delete by issuing the following command:

```
kubectl describe node node
```

- b. Delete the pods by issuing the following command:

```
kubectl delete pods pods_list
```

- c. Verify that the deleted pods are placed on another node by issuing the following command:

```
kubectl get pods -o wide
```

- d. Verify that the original node now has the capacity for the large pod that needs to be deployed by issuing the following command:

```
kubectl describe node node
```

## Results

Space is freed up on the specific node. Kubernetes finds this space on the node for Cassandra and deploys Cassandra on it. Continue with your Cloud App Management deployment.

## CEM datalayer pod repeatedly tries to restart and fails to start

### Problem

The CEM datalayer pod continuously tries to restart. After many attempts, the pod fails to start.

### Symptom

This issue is happening because the CEM datalayer pod has too many job files. As a result, the CEM datalayer pod is being killed by the liveness probe before it starts up.

### Solution

1. Increase the value of the **initialDelaySeconds** liveness probe from the default value of 120 to a much higher value, for example; 500. Enter the following these kubectl commands:

```
kubectl patch sts RELEASE-ibm-cem-datalayer -p
'{"spec":{"template":{"spec":{"containers":[{"name":"datalayer","livenessProbe":
{"initialDelaySeconds":500}}]}}}}'
```

where *RELEASE* is the helm release.

2. Verify the value of **initialDelaySeconds** is updated:

```
kubectl get sts RELEASE-ibm-cem-datalayer -o
jsonpath='{.spec.template.spec.containers[0].livenessProbe.initialDelaySeconds}'
```

The CEM datalayer pod starts successfully.

## The post-install-setup.sh script runs with the timeout error on Red Hat OpenShift environment

When you install IBM Cloud App Management on Red Hat OpenShift, you need to run the post-install-setup.sh script to complete the necessary administrative tasks. You might see the `Connection timed out` error if the Cloud Event Management pod and IBM Cloud Private pod use different namespaces.

### Problem

On some Red Hat OpenShift environments, services across different namespaces can't connect to each other. For example, the Cloud Event Management pod in the `icam` namespace can't connect to the IBM Cloud Private pod in the `kube-system` namespace. In this case, the post-install-setup.sh script fails with the `Connection timed out` error.

### Symptoms

You might see the following error after you run the **post-install-setup.sh** script to install IBM Cloud App Management on Red Hat OpenShift. For more information about the command, see step [“18”](#) on

page 147 of the Installing IBM Cloud App Management on Red Hat Open Shift topic.

```
Done.
curl: (7) Failed to connect to icp-management-ingress.kube-system port 443: Connection timed out
command terminated with exit code 1
ERROR 1, exiting.
2019.11.18 11:44:30 uxvnrwg001a6574: (/opt/ibm/ibm-cloud-management/ibm-cloud-appmgmt-prod)
# ./ibm_cloud_pak/pak_extensions/post-install-setup.sh --releaseName ibmcloudappmgmt --namespace edmed-esmcam-dev --instanceName ibmcloudappmgmt --advanced
Identified release name from helm: "ibmcloudappmgmt"
Mon Nov 18 11:46:51 EST 2019 Checking if all pods have become ready. This could take ~10 minutes after the initial helm install.
Mon Nov 18 11:46:55 EST 2019 All 51 pods in release ibmcloudappmgmt and namespace edmed-esmcam-dev are ready.
Mon Nov 18 11:46:55 EST 2019 Registering OIDC
Mon Nov 18 11:46:56 EST 2019 ICP secret platform-oidc-credentials is found
Mon Nov 18 11:46:58 EST 2019 ibmcloudappmgmt-ibm-cem-cem-users-59d8d79d4c-k2t79 is found
Registering IBM Cloud Event Management identity ...

Checking registration...
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload  Total   Spent    Left   Speed
  0     0     0     0     0     0     0     0    0  0  --:--:--  0:02:07  --:--:--    0curl: (7) Failed to connect to icp-management-ingress.kube-system port 443:
Connection timed out

Done.
command terminated with exit code 1
```

## Solution

Currently, the workaround is to join the network between the Cloud Event Management project and the IBM Cloud Private project. After the network is joined, both projects can access the pods and services of each other. For more information about how to join the network, see [Managing Networking](#).

## Results

You can run the **post-install-setup.sh** script successfully.

## Deleting the IBM Cloud Private service instance after the server is uninstalled

You must delete the IBM Cloud Private service instance before you uninstall the Cloud App Management server. If you did not delete the service instance before you uninstalled, complete the following steps to delete the service instance now.

## Problem

You cannot clean up the service instance that was created and tied to the Cloud App Management server when it was installed.

## Symptoms

You try to delete the IBM Cloud Private service instance but you cannot delete it because the Cloud App Management server is already uninstalled.

## Solution

To delete one or more service instances after you uninstall the Cloud App Management server, complete the following steps:

1. Edit the service instance properties in an editor. The following command opens the properties with the vi editor.

```
kubectl edit serviceinstance my_serviceInstanceName --namespace default
```

Where *my\_serviceInstanceName* is the name of the service instance. For example, *ibmcloudappmgmt* is the default service instance.

2. Delete the final line to detach the service instance from the IBM Cloud Private service catalog. By removing this line, you are preventing the service catalog from trying to delete the service instance which was associated with a Cloud App Management cluster service broker that was previously uninstalled.

```
kubernetes-incubator/service-catalog
```

3. Save the changes.

4. Delete the service instance by issuing the following command:

```
kubectl delete serviceinstance my_serviceInstanceName --namespace default
```

### Results

The service instance (s) is deleted.

## Troubleshooting agents

Troubleshoot IBM Cloud App Management agent installation and configuration issues.

Learn how to isolate and resolve problems with Cloud App Management. Verify that your issues are not related to operating system requirements, such as disk, memory, and CPU capacities. For more information about system requirements for Cloud App Management, see [“System requirements” on page 75](#). To get support, see [Support](#) in the product documentation.

## Db2 Monitoring

You may find more details about Db2 monitoring known issues.

### Cloud APM Server Db2 agent does not restart automatically on RHEL platform

#### Problem

The Cloud APM server Db2 agent does not restart automatically in scenarios listed below:

1. During TEMA patch installation
2. Switching agent connection from Cloud APM server to Cloud App Management server on RHEL platform.
3. Switching agent connection from Cloud App Management server to Cloud APM server server switch on RHEL platform.

#### Symptom

1. TEMA patch installation:  
Error as:

```
KCI1387W Failed to start instance of the ud component
```

2. Switching Db2 agent from Cloud APM Server to Cloud App Management:  
Error message :

```
Failure: Agent failed to start. Check the agent start log.
```

3. Switching Db2 agent from Cloud App Management to Cloud APM Server:  
Error message :

```
Failure: Agent failed to start. Check the agent start log.
```

#### Cause

The command used in the IBM Cloud App Management script fails to start the Db2 agent on RHEL platform.

#### Solution

1. Login with Db2 instance owner or user.
2. Execute the given command to start the agent automatically: **<agent\_install\_dir>/bin/db2-agent.sh start <instancename>**  
where **<instancename>** is the name of Db2 instance being monitored.

## Cloud APM Db2 agent configuration fails in case TEMA patch is applied after agent configuration

### Problem

The Cloud APM Db2 agent configuration fails in case TEMA patch is applied after agent configuration.

### Symptom

The terminal shows the following messages:

```
*****
KCIIN0524E Error attempting to build a merge file
java.io.FileNotFoundException: : /opt/IBM/ITM/tmp/.ud.rc (Permission denied)
*****
KCIIN0230E Unable to prompt for input...
```

### Solution

Manually grant write permission to db2 admin group on /CandleHome/tmp/.ud.rc.

OR

Remove /CandleHome/tmp/.ud.rc and configure the agent.

## Watchdog does not work for V6 TEMA patch, V8 TEMA patch and the Cloud APM Db2 agent

### Problem

Watchdog does not work for V6 TEMA patch, V8 TEMA patch and the Cloud APM Db2 agent.

### Symptom

The Watchdog functionality fails to restart agent automatically after system restart or after agent is stopped abruptly.

### Solution

In case of V6 TEMA patch, start the agent manually using **itmcmd** start agent command.

In case of V8 TEMA patch and the Cloud APM agent, start the agent manually using **db2-agent.sh** start agent command.

## Microsoft Active Directory Monitoring

You may find more details about Microsoft Active Directory monitoring known issues.

### Agent upgrade requires reinstall

#### Problem

Upgrade of the Microsoft Active Directory agent does not work.

#### Symptom

The newly added widgets are not visible on the Cloud App Management console.

#### Cause

The registry entry of the variable `IRA_CUSTOM_METADATA_LOCATION` used to upload the tarball was not updated with latest version after upgrade.

#### Solution

Uninstall the existing agent and install the latest agent.

## Label of Saturation in Golden Signals is not displayed

### Problem

Label is not displayed for Saturation in Golden Signals.

### Symptom

Label is not displayed accordingly for Saturation in Golden Signals.

### Cause

Not known.

### Solution

Not available.

## PostgreSQL Monitoring

You may find here more details about PostgreSQL Monitoring known issues.

### No data is displayed in all of the widgets for PostgreSQL Instance resource and Database resource.

#### Problem

No data is displayed in all of the widgets for PostgreSQL Instance resource and Database resource.

#### Symptom 1

Following exception occurs in logs:

```
<CANDLE HOME>/logs/kpn_JDBC_<instance_name>_trace.log java.io.IOException: Connection to 10.46.44.18:5432 refused. Check that the hostname and port are correct and that the postmaster is accepting TCP/IP connections.
```

#### Cause

Agent is not able to connect to remote PostgreSQL database.

#### Solution

Refer the steps given to resolve this issue:

1. Open `postgresql.conf` file located at `/var/lib/pgsql/12/data`
2. Update the **listen\_addresses** parameter to accept the connection from remote host. For example **listen\_addresses = '\*'**
3. Restart the PostgreSQL database server
4. Restart the PostgreSQL agent

#### Symptom 2

Following exception occurs in logs:

```
<CANDLE HOME>/logs/kpn_JDBC_<instance_name>_trace.log java.io.IOException: FATAL: no pg_hba.conf entry for host "<TEMA IP>", user "postgres", database "ibmdb", SSL off
```

#### Cause

Agent is not able to connect to remote PostgreSQL database.

## **Solution**

Refer the given steps to resolve this issue:

1. Open `pg_hba.conf` file located at `/var/lib/pgsql/12/data`
2. Update the IPv6 local connections section to accept the connection from remote host. For example, `host all all 0.0.0.0/0 md5`
3. Restart the PostgreSQL database server
4. Restart PostgreSQL agent

## **Microsoft .NET Monitoring**

You may find more details about Microsoft .NET monitoring known issues.

### **Label is not displayed for Errors in Golden Signal widget for .NET agent**

#### **Problem**

Label is not displayed for Errors in the Golden Signal widget for Microsoft .NET agent.

#### **Symptom**

N/A

#### **Cause**

Not known

#### **Solution**

Not known

## **Troubleshooting data collectors**

Troubleshoot IBM Cloud App Management data collector issues.

Learn how to isolate and resolve problems with Cloud App Management. Verify that your issues are not related to operating system requirements, such as disk, memory, and CPU capacities. For more information about system requirements for Cloud App Management, see [“System requirements” on page 75](#). To get support, see [Support](#) in the product documentation.

### **Fails to trigger the threshold of Request Rate after upgrading Node.js data collector**

After upgrading the Node.js data collector from Cloud App Management 2018.4.1 to 2019.2.0, the previous threshold that was created for Request Rate cannot be triggered in the new version of Node.js data collector.

#### **Problem**

After upgrading the Node.js data collector from Cloud App Management 2018.4.1 to 2019.2.0, the old threshold for metadata Request Rate cannot trigger any event in the new version of Node.js data collector.

#### **Symptom**

In Cloud App Management 2018.4.1, create a custom threshold for metadata Request Rate in the Node.js data collector, it can be triggered correctly. Then, upgrade the Node.js data collector to 2019.2.0, the threshold cannot trigger any event.

## Solution

The metadata Request Rate is modified when upgrading Node.js data collector to a new version. You need to create a new threshold for Request Rate. For instructions about how to create a threshold, see [“Managing thresholds” on page 755](#).

## Troubleshooting the console and dashboards

Troubleshoot IBM Cloud App Management console and dashboard issues.

Learn how to isolate and resolve problems with Cloud App Management. Verify that your issues are not related to operating system requirements, such as disk, memory, and CPU capacities. For more information about system requirements for Cloud App Management, see [“System requirements” on page 75](#). To get support, see [Support](#) in the product documentation.

## Users are unable to start the Application Monitoring UI. Users can see ERR\_TOO\_MANY\_REDIRECTS in browser.

### Symptom

Cannot start Application Monitoring link in the UI for IBM Cloud App Management V2019.4.0 installed with IBM Cloud Pak for Multicloud Management User sees ERR\_TOO\_MANY\_REDIRECTS in browser. Examination of URLs in browser tools reveals 302 error accessing the URL `https://(hostname)/cemui/launch/auth`.

### Solution

Apply the fix `2019.4.0-IBM-ICAM-SERVER-IF0001`. For details, see [IBM Cloud App Management 2019.4.0 2019.4.0-IBM-ICAM-SERVER-IF0001 Readme](#).

### Workaround

This workaround must be repeated for each user and session.

User can access the Incidents page and authenticate there. They can then browse to resources from the Incidents page.

After the credentials are cached, they can directly access the Application Monitoring UI again until they logout or a new session is established.

## IBM Cloud Pak for Multicloud Management email link not working

### Problem

A user for Cloud App Management with IBM Cloud Pak for Multicloud Management receives a welcome mail, email verification mail, or other notification email. The mail contains a link to IBM Cloud Pak for Multicloud Management. This link is broken.

### Cause

IBM Cloud Pak for Multicloud Management requires a cookie to be set before you are authorized to fetch stylesheets and scripts. This cookie gets set when a user navigates directly through the UI. If you launch directly into IBM Cloud Pak for Multicloud Management, the cookie is not set and the header is not properly styled.

### Solution

Log in to IBM Cloud Pak for Multicloud Management console first, and select **Event Management**.



## Incidents fail to load in All Incidents tab

### Problem

Incidents cannot be retrieved and an error message is displayed:

```
ERROR [ReadStage-8] 2018-10-04 13:35:20,161 StorageProxy.java:1906 - Scanned over 100001 tombstones during query 'SELECT * FROM datalayer.active_incidents WHERE subscription = de4bb7b4-d28a-449b-83cc-8cce205053e4 AND ORDER BY (incident_uuid DESC, active DESC) LIMIT 5000' (last scanned row partition key was ((de4bb7b4-d28a-449b-83cc-8cce205053e4), 8f180b50-c732-11e8-b3f4-efa234a5ac24, true)); query aborted
```

### Cause

By default, tombstones last for 10 days before they are cleaned up. A large number of events opening/closing over a 10 day period might cause this problem.

### Solution

The grace period for tombstones can be configured. The following command sets the time period to two days (172,800 seconds):

```
alter materialized view datalayer.active_incidents with gc_grace_seconds = 172800;
```

You can use `kubect1` (for Kubernetes) and `cqlsh` (for Cassandra) to submit this statement. See the following example:

```
$ kubect1 get pods
NAME                                READY   STATUS
RESTARTS   AGE
cemonicip-cassandra-0               1/1     Running
0         46h
cemonicip-couchdb-0                 1/1     Running
0         2m33s
cemonicip-ibm-cem-brokers-6564b45c47-68pbr 1/1     Running
1         46h
cemonicip-ibm-cem-cem-users-d4fdf67bc-l2k5b 1/1     Running
0         46h
cemonicip-ibm-cem-channelservices-65db684f96-x8m9q 1/1     Running
0         46h
cemonicip-ibm-cem-datalayer-0        1/1     Running
0         2m43s
cemonicip-ibm-cem-datalayer-cron-1556759700-xzzzl 0/1     Completed
0         7m2s
cemonicip-ibm-cem-event-analytics-ui-7cf786cc99-sqvw8 1/1     Running
0         46h
cemonicip-ibm-cem-eventpreprocessor-5d8d98dd54-s8nvl 1/1     Running
2         15h
cemonicip-ibm-cem-incidentprocessor-5bd964646-pkhg6 1/1     Running
0         15h
cemonicip-ibm-cem-integration-controller-7cb689fffd-kg94x 1/1     Running
1         46h
cemonicip-ibm-cem-normalizer-74dd497c94-zhlpn 1/1     Running
0         15h
cemonicip-ibm-cem-notificationprocessor-f9cd8d65d-dql141 1/1     Running
0         15h
cemonicip-ibm-cem-rba-as-545894d565-s6mn4 1/1     Running
0         15h
cemonicip-ibm-cem-rba-rbs-5cd7d44556-lp747 0/1     Running
0         15h
cemonicip-ibm-cem-scheduling-ui-5fbdd6649f-8svn6 1/1     Running
0         15h
cemonicip-kafka-0                   2/2     Running
0         2m34s
cemonicip-redis-sentinel-689c8764b5-9kldm 0/1     CrashLoopBackOff
11        15h
cemonicip-redis-server-6585cd65cc-277vn 0/1     CrashLoopBackOff
11        15h
cemonicip-zookeeper-0               1/1     Running
0         2m37s
<name>:deploy-ibm-cloud-private <name$> kubect1 exec -it cemonicip-cassandra-0 /bin/bash
cassandra@cemonicip-cassandra-0:/$ cd /opt/ibm/cassandra/bin
```

```

cassandra@cemonicp-cassandra-0:/opt/ibm/cassandra/bin$ ./cqlsh -u cassandra
Password:cassandra
Connected to apm_cassandra at 127.0.0.1:9042.
[cqlsh 5.0.1 | Cassandra 3.11.3 | CQL spec 3.4.4 | Native protocol v4]
Use HELP for help.
cassandra@cqlsh> alter materialized view datalayer.active_incidents with gc_grace_seconds =
172800;
cassandra@cqlsh> exit
cassandra@cemonicp-cassandra-0:/opt/ibm/cassandra/bin$ exit
exit

```

**Note:** If the grace period setting is changed to two days, for example, and a node is offline for two days, any data that was deleted more than two days before the node comes back online will be regenerated.

## IBM Cloud App Management resources view cannot display due to internalOAuthError message

An internalOAuthError message causes **Resources** view display issues.

### Problem

An internalOAuthError message is displayed when you try to access the **Resources** view from the **Resources** tab in the Cloud App Management UI.

### Symptoms

The **Resources** view cannot be displayed when you access it from the **Resources** tab. The problem relates to DNS resolution of the IBM Cloud Private proxy. The IBM Cloud Private proxy host name that is used for the Cloud App Management server installation is not known within Kubernetes, specifically the AMUI pod.

### Solution

1. Edit the IBM Cloud App Management UI deployment by issuing the following command:

```
kubectl edit deployment my_release_name-amui -n my_namespace
```

where *my\_release\_name* is the Cloud App Management release name. The default is `ibmcloudappgmt`. *my\_namespace* is default.

The deployment code is displayed in an editor.

2. Add the following four lines of code between the `dnsPolicy` and `restartPolicy` lines (near the end of file).

```

dnsPolicy: ClusterFirst
hostAliases:
- hostnames:
  - my_IBM_Cloud_Private_proxy_fully_qualified_domain_name
  ip: my_IBM_Cloud_Private_proxy_IP_address
restartPolicy: Always

```

3. Exit the editor with `:wq`.

### Results

A new AMUI deployment is automatically deployed. The new deployment contains an entry in its `/etc/hosts` file for the host name and IP address that is entered. As a result, the AMUI pod resolves and can access the proxy without adding a proxy to a DNS. You can now access the **Resources** view from the **Resources** tab.

## IBM Tivoli Monitoring threshold severities are incorrect in the UI

The severities for the IBM Tivoli Monitoring (ITM) thresholds are not displaying properly in the Cloud App Management UI.

### Problem

The threshold severities are not showing correctly in the Cloud App Management UI.

### Symptoms

For ITM default situations, the following severities are shown as **Indeterminate** in the Cloud App Management UI:

- HARMLESS
- INFORMATIONAL
- UNKNOWN

For the custom thresholds, the **MAJOR** severity is translated to **WARNING** on the ITM side, and this severity is displayed as **WARNING** in the incidents in the Cloud App Management UI.

### Solution

This is a known issue. No solution is available.

## Dashboards for Kubernetes resources are empty after Cloud App Management server upgrade

### Symptom

After you upgrade from one Cloud App Management server version to another, for example; from 2019.3.0 to 2019.4.0, the dashboards for runtime data collectors such as the Liberty data collector and Node.js data collector are not showing data. The UI pod log also shows errors connecting to Couch: `kubectl logs {amui pod} -c amuirest-service`. Even after you run the `kubectl delete pod my_pod_name` command, which restarts the runtime data collector to reregister the **Resources** dashboards, the dashboards are still empty.

### Cause

The connection with the CouchDB is lost during the page registration or initialization.

### Solution

1. Restart the data collector pod:

```
kubectl delete pod my_pod_name
```

2. Open the Resource dashboards for runtime data collectors to confirm that metrics are displayed.
3. If the dashboards are still not showing metrics, restart the UI pod: `kubectl delete pod amui`

## Troubleshooting general issues

Use this information to troubleshoot general IBM Cloud App Management problems.

Learn how to isolate and resolve problems with Cloud App Management. Verify that your issues are not related to operating system requirements, such as disk, memory, and CPU capacities. For more information about system requirements for Cloud App Management, see [“System requirements” on page 75](#). To get support, see [Support](#) in the product documentation.

## Events are being routed to the wrong account or team

### Symptoms

Cloud App Management events are being routed to the wrong account or team.

### Causes

The reason might be because a managed cluster is shared by a team used by the admin default account and another team.

### Resolving the problem

Ensure that managed clusters are only added to one team/cluster combination.

## Error message when deleting a service instance

When deleting a service instance using the Kubernetes command line tool, you might receive the following error:

```
**Error from server (BadRequest): the server rejected our request for an unknown reason**
```

This error message might be displayed even though the service instance has been successfully removed.

## Permission issue with Docker Version 18.03 with Ubuntu 16.04 LTS

If you use Docker Version 18.03 or higher with Ubuntu 16.04 LTS, containers that run as non-root might have permission issues. This issue appears to be due to a problem between the overlay storage driver and the kernel.

## Providing feedback

You can provide product feedback by opening a Request For Enhancement (RFE).

## Submitting an IBM Cloud App Management Request for Enhancement (RFE)

Request for Enhancement (RFE) is a way to submit an idea for a new feature or function for IBM Cloud App Management. Before you submit a new RFE request, search and view requests that are previously submitted. If your idea or a similar one is previously submitted, you can add comments to the existing request, which indicates your agreement with the idea. You can also vote for this idea. Product development teams might use this information to prioritize the development of new features.

### Procedure

1. Access the RFE community by using this URL: <https://www.ibm.com/developerworks/rfe/>  
A DeveloperWorks ID is required to use the RFE community. If you have an IBM external or client ID, you can use this same ID for the DeveloperWorks RFE site. If you do not have an ID, click **Register** on the upper right corner of the page. Complete the form and follow instructions to create an IBM ID.
2. Click **Sign In** to log in with your ID.
3. If it is your first time logging in to the RFE community, enter a display name and click **Continue**.  
You are redirected to the RFE community home page.
4. Click the **Submit** tab.
5. Complete the **Submit a request** form.  
All fields marked with an asterisk (\*) are mandatory.
  - a) In the **Product** field, enter IBM Cloud. A list of product names that begin with "IBM Cloud" are displayed. Select IBM Cloud App Management. The **Brand** and **Product family** fields automatically populate after the product name is selected.
  - b) Enter all relevant information.

It is important to enter as much useful information as possible. The IBM Cloud App Management product development team reviews this input and provides status updates about the decision they are making regarding this request.

- c) To add your vote to your RFE and add it to your watchlist, ensure the **Add vote** and **Add to To My watchlist** check boxes are selected.
  - d) Add any required attachments.
6. Click **Submit**.
  7. Review details on the request watchlist. You have 24 hours to confirm the entries that you submitted. You can use the **My stuff** tab. You can modify or delete the ticket item.

### What to do next

After the submitted RFE request is reviewed, the submitter is contacted by an IBM representative who uses the information that is provided in the RFE request.

## Support

If you have a problem with your IBM Cloud App Management software and you want to resolve it, the following topics describe how to collect logs if you have a server or agent issue, and how to contact IBM Support to help resolve the issue.

### Collecting the server logs for IBM Support

To troubleshoot server issues, and to enlist the Support team to help you, you must collect server logs for the IBM Support team.

#### Before you begin

Make sure to log in to Kubernetes before collecting the server logs.

#### About this task

To help you collect server logs, the `collectContainerLogs.sh` script is included in the Cloud App Management Passport Advantage Archive (PPA) installation image file.

You can get a description of the script and its options by issuing `./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/collectContainerLogs.sh --help`.

#### Procedure

1. If you are not actively working in the IBM Cloud Private cluster from the IBM Cloud Private CLI for a few hours, your authorization expires. Log in again by entering the following command:

```
cloudctl login -a https://my_cluster_name.icp:8443 --skip-ssl-validation
```

Where `my_cluster_name` is the IBM Cloud Private name defined for your cluster. The default value is `mycluster`.

2. Browse to the installation directory where you extracted the PPA file and run the following script: `install_dir/ibm_cloud_pak/pak_extensions/collectContainerLogs.sh`

A `tgz` file with a time stamp in the file name is generated in the `/tmp` directory, for example:

```
/tmp/diagnostic_data_20180610T024415Z.tgz.
```

3. Send the output file that is created to your IBM Support representative.

#### Results

The script gathers the following diagnostic information:

- Helm installation.

- Statistics that relate to the individual Kubernetes Pods.
- Artifacts that relate to the Cloud App Management server.
- Logs from the individual Pods.

## Collecting monitoring agent logs for IBM Support

Use the problem determination collection tool, *pdcollect*, to gather required logs and other problem determination information that is requested by IBM Support for monitoring agents. The PD collector tool is installed with each monitoring agent. It is applicable to both Cloud App Management agents and Cloud APM V8 agents.

### Before you begin

Root or administrator permission is required for the PD collector tool to collect system information from the monitoring agents. You can review the agent logs individually in the following folders:

- **Windows** [64-bit] `install_dir\TMAITM6_x64\logs`
- **Windows** [32-bit] `install_dir\TMAITM6\logs`
- **Linux** | **UNIX** `install_dir/logs`

where *install\_dir* is the agent installation directory. The default is as follows:

- **Windows** `C:\IBM\APM`
- **Linux** `/opt/ibm/apm/agent`
- **AIX** `/opt/ibm/apm/agent`

**Restriction:** It is only possible to run one instance of the *pdcollect* script.

**Note:** Most cases where you need to collect logs and system information from the monitoring agents also require collecting system information from the Cloud App Management server. When you collect the agent information, be sure to include the server information by running the `collectContainerLogs.sh` as documented here: [“Collecting the server logs for IBM Support” on page 1417](#).

### Procedure

To run the PD collector tool, complete the following steps:

1. On the command line, change to the agent directory:

- **Linux** | **UNIX** `install_dir/bin`
- **Windows** `install_dir\BIN`

2. Run the following command:

- **Linux** | **UNIX** `./pdcollect`
- **Windows** `pdcollect`

A compressed file with a time stamp in the file name is generated in the `tmp` or `Temp` directory, such as `/tmp/pdcollect-nc049021.tar.Z` or `/Temp/pdcollect-ADMIN_Tue06-12-2018.zip`.

3. Send the output files to your IBM Support representative.

## Contacting IBM Support

To help you troubleshoot issues when they arise, you can contact with IBM Support by opening support cases.

### Procedure

1. Go to the [IBM Support site](#).

- In the **Support Cases** section, click **Open a case**. You are prompted to enter your IBM ID and password and then you can see IBM Privacy Statement. Read the statement carefully and click **I consent** if you agree with the statement.

**Note:** If you don't agree with IBM Privacy Statement, click **cancel** to quit the login.

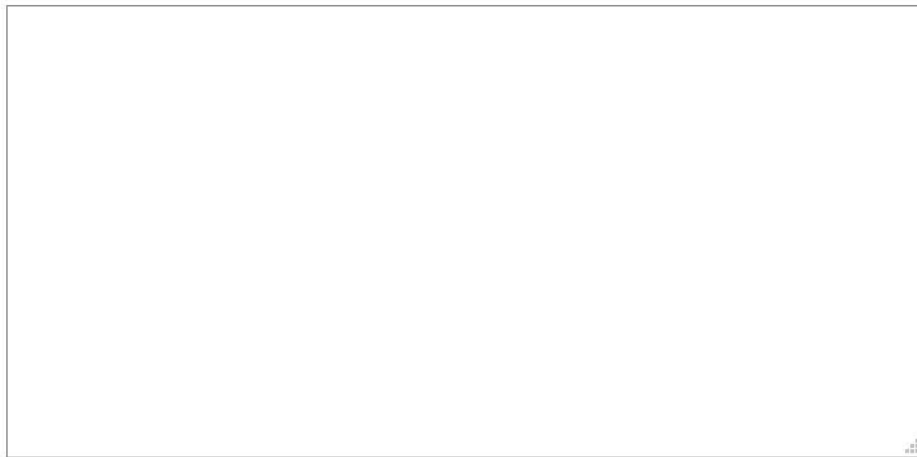
- In the **Open a case** page, enter the detailed information in the following fields:
  - In the **Title** field, enter your problem title.
  - In the **Product Manufacturer** field, enter IBM.
  - In the **Product** field, enter Cloud app Management.
  - In the **Severity** section, choose the case severity based on the business impact of the problem. Note if you set the case severity as 1, you need to have 24x7 availability to work with IBM support on this issue. See the following image for the severity details:

**\* Severity**

- 1 - Critical business impact (production or service is down)
- 2 - Significant impact (any system is down)
- 3 - Minor business impact
- 4 - Minimal impact (how-to questions, minor problems)

- In the **Description** field, enter the detailed problem that you met. A detailed description can help support understand your problem more accurately and thus provide solutions or answers more quickly.

**\* Description**



32,000 characters remaining

**Include the following information in**

Product version number (if available)

Problem

example: "we see error code/message"

Steps to reproduce

example: "tried restarting", "collected logs"

Suggestions/Answers you seek

example: "workaround", "help", "documentation"

Expected outcome

example: "reset password", "service restored"

- Select your preferred language, and choose whether you are willing to communicate in English when an agent who speaks your language is not available.
- Click **Submit case** to submit the case.

**Results**

IBM support can contact with you quickly by email or call after your case is submitted. You can get professional support in a timely manner.



---

## Chapter 23. Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. 2018.

## Trademarks

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

## Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

### Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

### Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

### Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

### Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

## IBM Online Privacy Statement

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth in the following paragraphs.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's user name for purposes of session management, authentication, and single sign-on configuration. These cookies can be disabled, but disabling them will also likely eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

---

## Chapter 24. Event filtering and summarization

Use the event filtering and summarization options that you set in the configuration (.conf) file to control how duplicate events are handled by the OS agent.

When a log is monitored, an event can display multiple times quickly. For example, this repeated logging can occur when the application that produces the log encounters an error and it logs this error continuously until the threshold is resolved. When this type of logging occurs, an excessive number of events are sent to the Performance Management infrastructure. The volume of events has a negative impact on performance.

**Note:** The event detection and summarization procedures are supported only on events that are sent to Performance Management. You cannot complete these procedures on events that are sent to OMNIBus by EIF.



## Chapter 25. Windows Event Log

The OS agent uses the .conf file to monitor events from the Windows Event Log.

The OS agent continues to use the WINEVENTLOGS configuration ( .conf) file option to monitor events from the Windows Event Log. The agent monitors a comma-separated list of event logs as shown in the following example:

```
WINEVENTLOGS=System,Security,Application
```

The OS agent also continues to use the WINEVENTLOGS=All setting. The All setting refers to the following standard event logs: Security, Application, System, Directory, Domain Name System (DNS), and File Replication Service (FRS) that come with Windows versions earlier than 2008. However, all the event logs on the system are not checked.

The UseNewEventLogAPI configuration file tag allows the event log (Windows Event Log 2008 or later) to access any new logs added by Microsoft, and any Windows event logs created by other applications or the user. The new logs are listed by the WINEVENTLOGS keyword.

In the following example, the UseNewEventLogAPI tag is set to y.

```
UseNewEventLogAPI=y  
WINEVENTLOGS=Microsoft-Windows-Hyper-V-Worker-Admin
```

In this example, the Microsoft-Windows-Hyper-V/Admin is monitored on a Windows system that has the Hyper-V role.

In the Windows Event Log, each event has the following fields in this order:

- Date in the following format: month, day, time, and year
- Event category as an integer
- Event Level
- Windows security ID. Any spaces in the Windows security ID are replaced by an underscore if SpaceReplacement=TRUE in the configuration ( .conf) file.

**Note:** SpaceReplacement=TRUE is the default if you set UseNewEventLogAPI to y in the ( .conf) file (designated that you are using the event log).

- Windows source. Any spaces in the Windows source are replaced by an underscore if SpaceReplacement=TRUE in the configuration ( .conf) file.
- Windows event log keywords. Any spaces in the Windows event log keywords are replaced by an underscore if SpaceReplacement=TRUE in the configuration ( .conf) file.

**Note:** The keyword field that is described here is new to the Windows 2008 version of Event Log. It did not exist in the previous Event Log, and so its presence prevents you from reusing your old Event Log format statements directly. They must be modified to account for this additional field.

- Windows event identifier
- Message text

For example, when an administrative user logs on to a Windows 2008 system, an event is generated in the Security log indicating the privileges that are assigned to the new user session:

```
Mar 22 13:58:35 2011 1 Information N/A Microsoft-Windows-  
Security-Auditing Audit_Success 4672 Special privileges assigned to new logon.  
S-1-5-21-586564200-1406810015-1408784414-500 Account Name:  
Administrator Account Domain: MOLDOVA Logon ID:  
0xc39cb8e Privileges: SeSecurityPrivilege  
SeBackupPrivilege SeRestorePrivilege  
SeTakeOwnershipPrivilege SeDebugPrivilege  
SeSystemEnvironmentPrivilege SeLoadDriverPrivilege  
SeImpersonatePrivilege
```

To capture all events that were created by the Microsoft-Windows-Security-Auditing event source, you write a format statement as shown here:

```
REGEX BaseAuditEvent
^([A-Z][a-z]{2} [0-9]{1,2} [0-9]{1,2}:[0-9]{2}:[0-9]{2} [0-9]{4}) [0-9] (\S+) (\S+) Microsoft-Windows-Security-Auditing (\S+) ([0-9]+) (.*)
timestamp $1
severity $2
login $3
eventsourc "Microsoft-Windows-Security-Auditing"
eventkeywords $4
eventid $5
msg $6
END
```

For the previous example event, the following example indicates the values that are assigned to slots:

```
timestamp=Mar 22 13:58:35 2011
severity=Information
login=N/A
eventsourc=Microsoft-Windows-Security-Auditing
eventid=4672
msg="Special privileges assigned to new logon.
S-1-5-21-586564200-1406810015-1408784414-500 Account Name:
Administrator Account Domain: MOLDOVA Logon ID:
0xc39cb8e Privileges: SeSecurityPrivilege
SeBackupPrivilege SeRestorePrivilege
SeTakeOwnershipPrivilege SeDebugPrivilege
SeSystemEnvironmentPrivilege SeLoadDriverPrivilege
SeImpersonatePrivilege
```

Because it is difficult to anticipate exactly what these events look like, a useful approach to writing your regular expressions is to capture the actual events in a file. Then, you can examine the file, choose the events that you want the agent to capture, and write regular expressions to match these events. To capture all events from your Windows Event Log, use the following steps:

1. Create a format file that contains only one pattern that does not match anything, as shown in the following example:

```
REGEX NoMatch
This doesn't match anything
END
```

2. Add the following setting to the configuration (.conf) file:

```
UnmatchLog=C:/temp/evlog.unmatch
```

3. Run the agent and capture some sample events.



# Chapter 26. Creating a Scripting REST API synthetic test

Use a scripting REST API test to test a sequence of REST APIs. Use a node . js script to test your sequenced REST APIs.

## Procedure

### Name and Description

1. Enter a meaningful name for your test in the **Name** field. Add a description of the purpose of your test to the **Description** field.

### Test type

2. Select Scripting REST API.

### Request

3. You can choose one of the following options from the **Upload options** list to provide the test script:

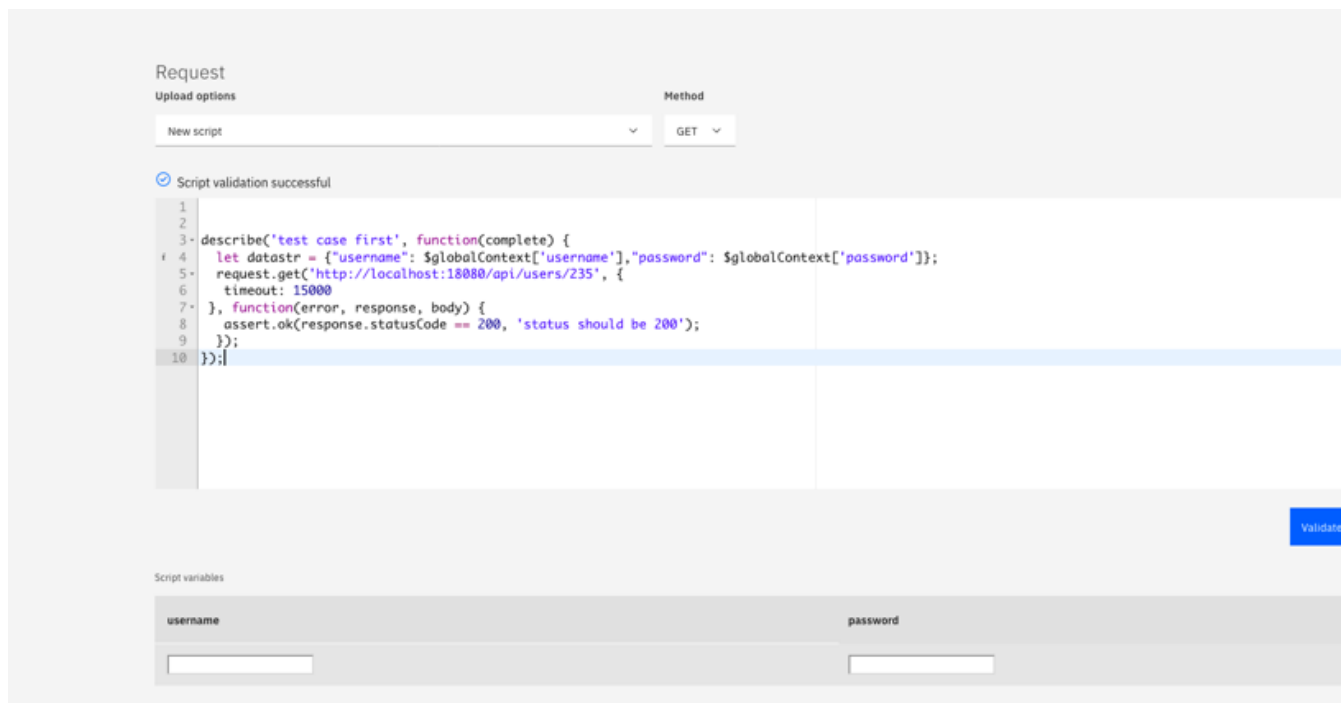
- **Upload a JS file**

Use the **Upload a JS file** option to upload a node . js test script file. For information about how to create a node . js script to test, see [“Create a REST API test case”](#) on page 629.

- **New script**

Use the **New script** option to create a new test script. Type your script in the script editor. Choose a method from the **Method** list. Then the template script for the method that you choose is added to the script editor. You only need to edit the template script to meet your goal, instead of typing the script manually.

If you use `$globalContext[VAR_NAME]` to pass variables in the script, you can see the **Script variables** section after clicking **Validate**. You can type the variable values in the **Script variables** section. The **Script variables** section is hidden by default.



**Note:** You can verify whether your script is grammatically right by clicking **Validate**. You can also download the script to your local computer by clicking **Download**.

- **Add a template**

This option works in the same way as the **New script** option and is designed to be removed from Cloud App Management console **V2020.1.0**.

#### Response validation

4. Configure the warning and critical events conditions for your synthetic test in the **Response Validation** section. You can see two conditions based on response time are provided to trigger events. By default, a response time over 5 seconds triggers a warning event and a response time over 10 seconds triggers a critical event. You can change the response time in the **Threshold Value** field or change the unit to milliseconds or seconds in the **Unit** field for each condition. Response times that exceed threshold values in your warning and critical conditions trigger events.

Further customization of warning and critical events can be done in the next configuration stage. For more information, see Event triggers later in this procedure.

For more detail in relation to event triggers default behavior and how event triggers function across multiple Synthetics PoP locations, see [“Event generation” on page 639](#).

#### Verify

5. Click **Verify** to determine whether your test request is valid. No response validation takes place during test verification. Your validated test is displayed in the verified test window. You can rename or delete your test in the verified test window. Click **Next**.

#### Review and Finish

6. Enter an Interval and Testing frequency.

##### **Interval**

Defines how often the test runs in minutes or hours.

##### **Testing frequency**

Determines whether your test runs from all locations simultaneously or from a different location at each interval. Select **Simultaneous** to run your test from all locations simultaneously, or select **Staggered** to run your test from a different selected location at each interval.

#### Locations

7. The **Locations** sections lists the Synthetics PoP that are installed. The first Synthetics PoP is selected by default. You can run your test from one or more synthetic pop servers.

Select the synthetic pop servers where you want your synthetic test to run. To create a new Location, see [“Installing Synthetics PoP” on page 620](#).

#### Script variables

8. If you introduced any variables in the node . js script, they are requested at this point.

#### Event triggers

9. By default a critical alert is triggered if a synthetic test playback fails (returns a code 400 or above).

To stop this behavior, set **Trigger an event if a failure is detected** to **Off**. To increase the number of failures allowed before a critical alert is triggered, change the value between **Trigger an event if the test fails** and **consecutive times** under the **Failure** section. The default number of consecutive failures is 0.

By default, a critical event is triggered if a synthetic test playback response time is >10 seconds. By default, a warning event is triggered if a synthetic test playback response time is >5 seconds.

To increase the number of slow response times that must occur before a critical or warning event is triggered, change the value between **Trigger an event if a threshold is breached** and **consecutive times** under the **Slow response threshold** section. The default number of slow response times is 0.

For more detail in relation to event triggers default behavior and how event triggers function across multiple Synthetics PoP locations, see [“Event generation” on page 639](#).

# Accessibility features

---

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

## Accessibility features

The web-based interface of IBM Cloud App Management is the Cloud App Management console. The console includes the following major accessibility features:

- Enables users to use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen. Consult the product documentation of the assistive technology for details on using those technologies with this product.
- Enables users to operate specific or equivalent features using only the keyboard.
- Communicates all information independently of color.<sup>2</sup>

The Cloud App Management console uses the latest W3C Standard, [WAI-ARIA 1.0 \(http://www.w3.org/TR/wai-aria/\)](http://www.w3.org/TR/wai-aria/), to ensure compliance with [US Section 508 \(http://www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards\)](http://www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards), and [Web Content Accessibility Guidelines \(WCAG\) 2.0 \(http://www.w3.org/TR/WCAG20/\)](http://www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader in combination with the latest web browser that is supported by this product.

The Cloud App Management console online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described at [IBM Knowledge Center release notes http://www.ibm.com/support/knowledgecenter/doc/kc\\_help.html#accessibility](http://www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

## Keyboard navigation

This product uses standard navigation keys.

## Interface information

The Cloud App Management console web user interface does not rely on cascading style sheets to render content properly and to provide a usable experience. However, the product documentation does rely on cascading style sheets. IBM Knowledge Center provides an equivalent way for low-vision users to use their custom display settings, including high-contrast mode. You can control font size by using the device or browser settings.

The Cloud App Management console web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

The Cloud App Management console user interface does not have content that flashes 2 - 55 times per second.

## Related accessibility information

In addition to standard IBM help desk and support websites, IBM has established a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service 800-IBM-3383 (800-426-3383) (within North America)

## IBM and accessibility

For more information about the commitment that IBM has to accessibility, see [IBM Accessibility \(www.ibm.com/able\)](http://www.ibm.com/able).

---

<sup>2</sup> Exceptions include some **Agent Configuration** pages of the Performance Management console.



## Notices

---

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. 2018.

## Trademarks

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

## **Terms and conditions for product documentation**

Permissions for the use of these publications are granted subject to the following terms and conditions.

### **Applicability**

These terms and conditions are in addition to any terms of use for the IBM website.

### **Personal use**

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

### **Commercial use**

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

### **Rights**

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

## **IBM Online Privacy Statement**

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth in the following paragraphs.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's user name for purposes of session management, authentication, and single sign-on configuration. These cookies can be disabled, but disabling them will also likely eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.





