

IBM Cloud App Management  
Version 2019 Release 3

*User's Guide*



**Note**

Before using this information and the product it supports, read the information in [Chapter 5, “Notices,”](#) on page 47.

This edition applies to version 2019.3.0 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. PDF documentation.....</b>	<b>11</b>
Documentation.....	11
Conventions used in the documentation.....	11
<b>Chapter 3. Product overview.....</b>	<b>13</b>
Offerings.....	14
User interface.....	16
Getting started: Manage dynamic application and infrastructure environments .....	17
Getting started: Collaborate to rapidly resolve problems .....	20
Getting Started: Proactively manage the health of your application environment – regardless of size .....	24
Getting started: Accelerate your transition to the cloud with DevOps.....	28
Getting started: Performing SRE functions.....	33
Agents and data collectors.....	35
Change history.....	37
Descriptions.....	38
<b>Chapter 4. Accessibility features.....</b>	<b>45</b>
Trademarks.....	48
Terms and conditions for product documentation.....	49
IBM Online Privacy Statement.....	49
<b>Chapter 6. Planning your deployment.....</b>	<b>51</b>
Product components.....	52
Part numbers.....	53
System requirements.....	57
Planning hardware and sizing .....	58
Planning ports.....	61
Scaling stateless and stateful services .....	63
Scaling Cassandra .....	64
Scaling up CouchDB.....	65
Scaling up Kafka brokers.....	66
Scaling up ZooKeeper.....	67
Scaling up Datalayer.....	68
<b>Chapter 7. Deploying IBM Cloud Private.....</b>	<b>69</b>
Preparing to install IBM Cloud Private.....	69
Installing an IBM Cloud Private Enterprise environment.....	70
IBM Cloud Private postinstallation manual tasks.....	72
<b>Chapter 8. Installing IBM Cloud App Management - the options.....</b>	<b>73</b>
Installing IBM Cloud App Management with IBM Cloud Pak for Multicloud Management.....	73
Onboarding LDAP users.....	73
Deploying the Cloud App Management server with IBM Cloud Pak for Multicloud Management.....	75
Deploying the ICAM klusterlet for IBM Multicloud Manager .....	83
Deploying the ICAM klusterlet for IBM Multicloud Manager without helm.....	85
Deploying agents and data collectors for IBM Multicloud Manager.....	89

Uninstalling Cloud App Management for IBM Multicloud Manager.....	90
Installing IBM Cloud App Management on Red Hat OpenShift.....	91
Installing IBM Cloud App Management on IBM Cloud Private .....	99
Downloading the Cloud App Management PPA file.....	100
Planning for a high availability installation.....	100
Certificates for HTTPS communications.....	104
Optimizing performance.....	107
Installing the Cloud App Management server.....	112
Creating your service instance.....	123
Starting the Cloud App Management UI.....	124
Uninstalling IBM Cloud App Management .....	125
Backing up and restoring .....	127

**Chapter 9. Deploying ICAM Agents..... 131**

Planning agent deployment.....	131
Downloading agents and data collectors from Passport Advantage.....	132
Configuring the downloaded images.....	132
Installing agents on UNIX systems.....	134
Preinstallation on AIX systems.....	134
Preinstallation on Solaris systems.....	135
Installing agents.....	136
Installing agents on Linux systems.....	138
Preinstallation on Linux systems.....	138
Installing agents.....	139
Installing agents on Windows systems.....	141
Preinstallation on Windows systems.....	142
Installing agents.....	142
Installing agents silently.....	144
Installing agents as a non-root user.....	145
Postinstallation tasks for the agents.....	146
Securing the agent installation files.....	148
Bypassing the prerequisite scanner.....	149
Uninstalling your agents.....	150
WebSphere Applications agent: Unconfiguring the data collector.....	152

**Chapter 10. Configuring your agents..... 161**

Common procedures.....	161
Using agent commands.....	162
Starting agents as a non-root user.....	166
Configuring agents as a non-root user.....	167
Disabling automatic agent start on AIX and Linux systems.....	168
Configuring Cisco Unified Computing System (UCS) monitoring.....	168
Creating a user and granting required permissions.....	169
Configuring the agent on Windows systems.....	169
Configuring the agent by using the silent response file.....	170
Configuring the agent on Linux systems.....	171
Configuration parameters for the agent.....	172
Configuration parameters for the data provider.....	172
Enabling SSL communication with Cisco UCS data sources.....	173
Increasing the Java heap size.....	174
Configuring DataPower monitoring.....	174
Configuring DataPower appliances.....	175
Configuring the DataPower agent.....	176
Configuring Db2 monitoring .....	178
Configuring the agent on Windows systems.....	180
Configuring the agent on Linux or UNIX systems.....	181
Configuring the agent by using the silent response file.....	182

Granting privileges for viewing Db2 agent metrics.....	184
Configuring local environment variables.....	185
Prerequisites for Remote Monitoring.....	187
Configuring Hadoop monitoring.....	189
Configuring the agent on Windows systems .....	190
Configuring the agent on Linux and AIX systems.....	192
Configuring the agent by using the silent response file.....	195
Granting permission to non-admin users.....	196
Configuring the dashboard for viewing Hadoop events.....	197
Configuring HTTP Server agent monitoring.....	197
HTTP Server agent code samples.....	199
<b>Configuring IBM Integration Bus monitoring.....</b>	<b>200</b>
Configuring the IBM Integration Bus agent.....	201
Configuring IBM Integration Bus for data enablement.....	204
Disabling snapshot data collection for the agent.....	208
Configuring InfoSphere DataStage monitoring.....	209
Configuring the agent on Windows systems.....	209
Configuring the agent on Linux systems.....	209
Configuring environment variables.....	210
Configuring the agent by using the silent response file.....	210
Configuration parameters of the agent.....	211
Configuring JBoss monitoring.....	212
Add a JBoss Server Management User.....	213
Enable JMX MBean Server Connections.....	214
Enabling Web/HTTP Statistic Collection.....	215
Configuration Parameters for the JBoss agent.....	217
Configuring the agent on Windows systems.....	218
Configuring the agent by responding to prompts.....	218
Configuring the agent by using the silent response file.....	219
Configuring Linux KVM monitoring.....	220
Creating a user and granting required permissions.....	221
Configuring protocols.....	221
Configuring a connection to the RHEVM server.....	225
Configuring a connection to the RHEVH server.....	227
Configuration parameters to connect to the RHEVM server.....	227
Configuration parameters to connect to the RHEVH server.....	229
Configuring Microsoft .NET monitoring.....	231
Permissions to run an agent by using a local or domain account.....	231
Configuring the agent on Windows systems.....	232
Changing the user account.....	232
Configuring Microsoft Cluster Server monitoring.....	233
Creating a generic service cluster resource on Windows Server 2008, 2012, 2016, and 2019 systems .....	233
Configuring the agent by using the silent response file.....	234
Changing the user account.....	234
Configuring Microsoft Exchange Server monitoring.....	235
Creating users.....	235
Assigning administrator rights to the Exchange Server user.....	238
Making the Exchange Server user a local administrator.....	238
Configuring the Exchange Server for reachability.....	240
Configuring the agent to run under the domain user.....	241
Configuring the agent locally.....	241
Configuring the agent by using the silent response file.....	245
Configuring local environment variables for the agent.....	246
Configuring Microsoft Hyper-V monitoring.....	246
Providing Local Security Policy for running Monitoring Agent for Microsoft Hyper-V Server on Windows by a Non-Administrator user.....	247
Granting Local Security Policy permissions.....	248

Modifying DCOM permissions.....	249
Adding a non-administrator user in the Hyper-V administrator users group.....	249
Adding a non-administrator user in the Performance Monitor users group.....	250
Configuring Microsoft IIS monitoring.....	250
Configuring the agent on Windows systems.....	251
Configuring the agent by using the silent response file.....	251
Changing the user account.....	252
Configuring Microsoft Office 365 monitoring.....	252
Enabling remote execution policy for PowerShell.....	253
Verifying reachability of configured users.....	254
Configuring the agent on Windows systems.....	254
Configuring the agent by using the silent response file.....	255
Changing the user account.....	256
Monitoring the Skype's quality of service.....	256
Configuring local environment variables.....	257
Configuring Microsoft SharePoint Server monitoring.....	259
Changing the user account.....	259
Running Monitoring Agent for Microsoft SharePoint Server by a non-admin user.....	260
Local Security Policy permissions.....	261
Configuring Microsoft SQL Server monitoring.....	261
Creating a user and granting permissions.....	262
Running as a non-administrator user.....	266
Selecting the databases for monitoring.....	269
Configuring local environment variables.....	278
Configuring the agent by using the silent response file.....	284
Running the agent in a cluster environment.....	285
Configuring the agent by using the cluster utility.....	287
Configuring multiple collations for ERRORLOG file.....	288
Configuring historical job count for monitoring.....	289
Configuring MongoDB monitoring.....	290
Configuring the agent with default settings.....	292
Configuring the agent by responding to prompts.....	292
Configuring the agent by using the silent response file.....	293
Configuring MySQL monitoring.....	295
Configuring the agent on Windows systems.....	295
Configuring the agent on Linux systems.....	296
Configuring the agent by using the silent response file.....	296
Configuring NetApp Storage monitoring.....	298
Downloading and installing the NetApp Manageability SDK JAR file.....	298
Configuring the agent on Windows systems.....	299
Configuring the agent by using the silent response file.....	299
Configuring the agent on Linux systems.....	300
Configuration parameters for the data provider.....	301
Configuration parameters for the OnCommand Unified Manager.....	302
Configuration parameters for the OnCommand API Service.....	303
Configuring Oracle Database monitoring.....	303
Configuring the agent on Windows systems.....	305
Configuring the agent by responding to prompts.....	308
Configuring the agent by using the silent response file.....	312
Granting privileges to the Oracle Database agent user.....	314
Configuring OS monitoring.....	317
Run OS agents as a non-root user.....	317
Configure OS agent log file monitoring.....	319
Configuring Linux OS Agent file system data collection.....	339
Configuring PostgreSQL monitoring.....	339
Configuring the agent on Windows systems.....	340
Configuring the agent on Linux systems.....	341
Configuring the agent by using the silent response file.....	342

Configuring RabbitMQ monitoring .....	343
Configuring the agent on Windows systems.....	343
Configuring the agent on Linux systems.....	343
Configuring the agent by using the silent response file.....	344
Configuration parameters for the agent.....	345
Configuring SAP monitoring.....	345
Configuring the agent on Windows systems.....	346
Configuring the agent on Linux or AIX systems.....	347
Configuring the agent by using the silent response file.....	348
Configuration parameters of the agent.....	348
Importing the ABAP transport on the SAP system.....	351
Verifying agent configuration.....	358
Advanced installation and configuration of the SAP agent.....	362
Configuring SAP HANA Database monitoring.....	374
Configuring the agent on Windows systems.....	375
Configuring the agent on Linux and AIX systems.....	376
Configuring the agent by using the silent response file.....	376
Configuring SAP NetWeaver Java Stack monitoring.....	377
Configuring the agent on Windows systems.....	378
Configuring the agent on Linux or AIX systems.....	378
Configuring the agent by using the silent response file.....	379
Configuration parameters of the agent.....	380
Configuring Skype for Business Server monitoring.....	380
Permissions and access rights for a non-administrator user.....	381
Configuring the agent on Windows systems.....	382
Configuring the agent by using the silent response file.....	383
Changing the user account.....	383
Configuration parameters for the agent.....	384
Configuring Tomcat Monitoring .....	385
Configuring Tomcat agent with the default settings.....	386
Configuring Tomcat agent on Windows systems.....	387
Configuring Tomcat agent on Linux systems.....	388
Configuring Tomcat agent using silent response file.....	388
Configuring VMware VI monitoring.....	389
Sizing and planning the VMware VI agent deployment.....	390
Enabling SSL communication with VMware VI data sources.....	391
Configuring the agent on Windows systems.....	392
Configuring the agent by responding to prompts.....	393
Configuring the agent by using the silent response file.....	393
Configuration parameters for the data source.....	394
Configuration parameters for the data provider.....	395
Increasing the Java heap size.....	396
Configuring WebSphere Applications monitoring.....	397
Fast track: Configuring the data collector for WebSphere Applications agent.....	397
Checking user access requirements.....	399
Configuring the data collector with the simple configuration utility.....	400
Configuring or reconfiguring the data collector with full configuration utilities.....	401
Manually configure the data collector if the configuration utilities fail.....	416
Restoring the application server configuration from a backup.....	419
Configuring WebSphere Infrastructure Manager monitoring.....	421
Configuring WebSphere MQ monitoring.....	421
Authorizing the user IDs to run the agent.....	422
Configuring IBM MQ (WebSphere MQ) for data enablement.....	423
Configuring the IBM MQ (formerly WebSphere MQ) agent.....	424
<b>Chapter 11. Deploying ICAM Data Collectors.....</b>	<b>427</b>
Kubernetes data collector.....	429

Configuring Kubernetes monitoring.....	429
Configuring Kubernetes monitoring without Ansible.....	434
Configuring Kubernetes monitoring without Helm.....	437
Validating the Kubernetes installation logs.....	440
Kubernetes metrics for thresholds.....	441
Uninstalling the Kubernetes data collector.....	445
Kubernetes data collector issues.....	446
Runtime data collectors.....	448
Common topics.....	448
Configuring Node.js application monitoring.....	450
Configuring Liberty application monitoring.....	456
Configuring J2SE application monitoring.....	465
Configuring Python application monitoring.....	470
Synthetics PoP.....	480
Installing a Synthetics PoP.....	480
Creating a REST API test.....	483
Creating a Scripting REST API synthetic test.....	486
Creating a web page synthetic test.....	490
Creating a Selenium script test.....	491
Event generation.....	496
Viewing Synthetic test results.....	497
<b>Chapter 12. Unified Agent.....</b>	<b>499</b>
Overview of Unified Agent.....	499
Preparing the deployment of Unified Agent.....	499
Enabling the NGINX monitoring interface.....	502
Authorizing the plug-ins to access Kubernetes resources.....	505
Deploying the Unified Agent.....	506
Configuring Jaeger and Zipkin in Unified Agent.....	508
Configuring NGINX monitoring in Unified Agent.....	509
Configuring Redis monitoring in Unified Agent.....	509
Configuring IBM API Connect monitoring in Unified Agent.....	510
Configuring IBM App Connect Enterprise monitoring in Unified Agent.....	510
Configuring IBM MQ monitoring in Unified Agent.....	511
Updating the Unified Agent configuration.....	512
Uninstalling the Unified Agent.....	514
<b>Chapter 13. Integrating with other products.....</b>	<b>517</b>
Configuring monitoring data sources.....	517
Integrating with IBM Tivoli Monitoring agents.....	517
Integrating with Cloud APM, Private agents.....	530
Configuring incoming event sources.....	534
Creating custom event sources with JSON.....	534
Configuring Amazon Web Services (AWS) as an event source.....	535
Configuring AppDynamics as an event source.....	536
Configuring Datadog as an event source.....	538
Configuring Dynatrace as an event source.....	539
Configuring Elasticsearch as an event source.....	539
Configuring Jenkins as an event source.....	542
Configuring Logstash as an event source.....	543
Configuring Microsoft Azure as an event source.....	544
Configuring Nagios XI as an event source.....	545
Configuring Netcool/OMNIBus as an event source.....	547
Configuring New Relic as an event source.....	547
Configuring Pingdom as an event source.....	549
Configuring Prometheus as an event source.....	550
Configuring SolarWinds Orion as an event source.....	554



Configuring Splunk Enterprise as an event source.....	555
Configuring IBM UrbanCode Deploy as an event source.....	557
Configuring outgoing event destinations.....	559
Sending incident details to Alert Notification.....	559
Configuring Netcool/OMNIbus as an outgoing incident source.....	560
Sending incident details to Slack channels.....	562
Sending incident details using outgoing webhooks.....	563
Sending incident details to Microsoft Teams.....	564
Sending incident details to ServiceNow.....	565
Sending incident details to GitHub.....	566
Sending incident details to Stride.....	567
Sending incident details to Watson Workspace.....	567
<b>Chapter 14. Administering.....</b>	<b>569</b>
Setting up event policies.....	569
Events and incidents.....	570
Example: Changing event information through enrichment.....	573
Example: Adding to event information through enrichment.....	574
Example: Correlating events through enrichment.....	575
Creating lookup tables.....	576
Example: Enriching event information using lookup tables.....	577
Enriching event information by adding custom attributes .....	579
Example: Suppressing temporary high memory usage warnings.....	580
Example: Preventing certain events from forming incidents.....	581
Example: Setting runbook for disk full events .....	582
Example: Detecting flapping events.....	583
Managing incident policies.....	584
Example: Assigning top priority incidents automatically.....	585
Example: Escalating incidents automatically.....	586
Example: Setting incident priority.....	587
Managing incidents.....	589
Understanding your incidents at a glance.....	591
Understanding the mean time to resolve incidents.....	594
Starting to work with incidents.....	595
Resolving incidents with runbooks.....	596
Managing thresholds.....	597
Managing resource groups.....	600
Setting up users and groups.....	601
Logging in to the Cloud App Management UI to authenticate as a user.....	601
Activating notifications via SMS and voice.....	602
Defining users and groups.....	602
Setting up groups.....	604
Roles.....	604
Using the REST API.....	605
Managing runbooks.....	606
Data retention and summarization.....	606
Configuring summarization.....	607
<b>Chapter 15. Monitoring.....</b>	<b>611</b>
Viewing your managed resources.....	611
Resource dashboard.....	612
Service dependencies topology view.....	617
Monitoring the status of your Tivoli Monitoring data providers.....	622
Viewing Tivoli Monitoring data providers.....	622
<b>Chapter 16. Upgrading.....</b>	<b>625</b>
Upgrading your server.....	625

Upgrade to a high-availability environment.....	630
Upgrading your agents.....	634
Upgrading your ICAM Agents.....	635
Upgrading your IBM Tivoli Monitoring agents.....	636
Upgrading your Cloud APM agents.....	638
Upgrading your data collectors.....	639
Upgrading the Synthetics PoP server.....	640

## **Chapter 17. Troubleshooting (known issues, resolutions, and limitations) and support.....641**

Submitting an IBM Cloud App Management Request for Enhancement (RFE).....	641
IBM Multicloud Manager hub cluster fails to import a cluster with Cloud App Management server installed .....	642
IBM Cloud Pak for Multicloud Management email link not working.....	643
<b>Events are being routed to the wrong account or team.....</b>	<b>643</b>
Users accounts cannot be edited or saved.....	643
Incidents fail to load in All Incidents tab.....	643
Log Profile Event data not showing correctly.....	645
502 Bad Gateway error on IBM Cloud Private .....	645
Multiple service instances are not supported in same browser.....	646
No agent metrics after Kafka restarts.....	646
Error message when deleting a service instance.....	646
PPA installation file fails to load into the IBM Cloud Private repository with Docker authentication error.....	647
How do I reload support files to solve agent or environment problems .....	648
Cloud App Management installation fails with InvalidImageName error.....	649
Freeing up resources for large Pod scheduling.....	650
When scaling up ibm-redis, the ibm-redis-server pods crash.....	651
Resources view cannot display due to internalOAuthError message.....	652
Common deployment errors.....	653
Error during console login.....	654
IBM Tivoli Monitoring (ITM) threshold severities are incorrect in the UI.....	656
Deleting the IBM Cloud Private service instance after the Cloud App Management server is uninstalled.....	657
<b>V6 WebSphere Applications agent can't start after applying IF0004 agent patch on AIX 7.1....</b>	<b>657</b>
Cloud APM Server Db2 agent does not restart automatically on RHEL platform.....	658
Cloud APM Db2 agent configuration fails in case TEMA patch is applied after agent configuration....	659
Fresh installed Cloud APM Db2 agent on Windows fails to collect data.....	659
Permission error on the ud . env file when starting the second instance of APM Db2 agent with TEMA patch.....	660
Watchdog does not work for V6 TEMA patch, V8 TEMA patch and the Cloud APM Db2 agent.....	660
Fails to trigger the threshold of Request Rate after upgrading Node.js data collector.....	661
Dashboards for Kubernetes Resources are empty after Cloud App Management server upgrade.....	661
CEM datalayer pod repeatedly tries to restart and fails to start.....	662
Permission issue with Docker Version 18.03 with Ubuntu 16.04 LTS.....	662
Support.....	662
Collecting the server logs for IBM Support.....	662
Collecting monitoring agent logs for IBM Support.....	663
Contacting IBM Support.....	664

## **Accessibility features..... 671**

## **Notices.....673**

Trademarks.....	674
Terms and conditions for product documentation.....	675
IBM Online Privacy Statement.....	675

---

# 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 37.

## 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 168.

#### 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 212.

#### 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 233

#### 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 259.

#### 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 290.

#### 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 295.

### **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 339.

### **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 377.

## **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 Multicloud Manager”](#) on page 89.

### **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 12, “Unified Agent,”](#) on page 499.

### **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 [“Configuring digital experience monitoring \(DEM\) for Liberty applications”](#) on page 462.
- 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 606.

### 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 606.

## 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 600](#).

### 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 597](#).

### 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 617](#).

### 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 622](#).

## 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 12, “Unified Agent,” on page 499](#).

### 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 235](#).

#### 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 252](#).

### **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 on the IBM Cloud Application Management UI.

For information about configuring the agent after installation, see [“Configuring NetApp Storage monitoring”](#) on page 298.

## **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 209.

#### **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 189.

#### **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 380.

### **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 476.

### **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 [“Installing IBM Cloud App Management on Red Hat OpenShift”](#) on page 91.

### **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 73.

### **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 100.

### **Backup and restore**

Backup and restore your workload. For more information, see [“Backing up and restoring”](#) on page 127.

### **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 [“Configuring digital experience monitoring \(DEM\) for Liberty applications”](#) on page 462.

### **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 374.

### 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 421.

### 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 470.

## 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 200.

### 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 [“Configuring digital experience monitoring \(DEM\) for Liberty applications ” on page 462.](#)

For information about the agent version in each release or refresh, see [“Change history” on page 37.](#)

## 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 139](#) section of [“Preinstallation on Linux systems” on page 138.](#)

#### 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 149.](#)

## 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 231.](#)

#### 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 250.

### **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 261.

### **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 345.

### **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 385

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

### **New agents**

#### **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 212.

#### **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 246.

#### **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 220.

### **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 389.

## **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 465.

## **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 597.

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 612 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 480.

## **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 37.

---

## Chapter 2. 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.

### IBM Redbooks®

The IBM Redbooks website contains Redbooks publications, Redpapers, and Redbooks technotes that provide information about products from platform and solution perspectives.

## 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 3. Product overview

IBM Cloud App Management offers a comprehensive management solution for development, operations, and site reliability engineering (SRE) teams. It's a cloud native, container-based platform for managing your applications in a hybrid cloud environment.

What value does IBM Cloud App Management provide?

### Supports application modernization

Cloud App Management helps to minimize disruption to your monitoring operations as you modernize your applications and technologies.

Your organization might use IBM Tivoli Monitoring (V6) and IBM Cloud APM (V8) to monitor your environment. Legacy agents can run in dual mode with some agents reporting 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, 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

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.

### Unified agent

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. 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.

## Offerings

---

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.

The offerings include the following components.

### **Cloud App Management server**

For more information, see [“Downloading the Cloud App Management Passport Advantage Archive \(PPA\) file”](#) on page 100.

### **ICAM Agents**

For more information, see [“Descriptions”](#) on page 38.

### **ICAM Data Collectors**

Contains: Cloud data collectors, Node.js data collector, and Liberty data collector. For more information, see [“Descriptions”](#) on page 38.

### **Unified Agent**

Contains: Jaeger and Zipkin plug-in to receive OpenTracing workloads, NGINX and Redis plug-in to monitor NGINX and Redis workloads, IBM App Connect Enterprise and IBM MQ plug-ins to monitor IBM App Connect Enterprise and IBM MQ that are deployed in IBM Cloud Private environment, and IBM API Connect plug-in to monitor IBM API Connect. For more information, see [Chapter 12, “Unified Agent,”](#) on page 499.

### **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 53.



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>
Send incident details to Webhook	—	✓	<a href="#">How to</a>
Send incident details to Microsoft teams	—	✓	<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 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>• Jaeger and Zipkin plug-in</li> <li>• NGINX plug-in</li> <li>• Redis plug-in</li> <li>• IBM API Connect(APIC) plug-in</li> <li>• IBM App Connect Enterprise(ACE) plug-in</li> <li>• IBM MQ plug-in</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>

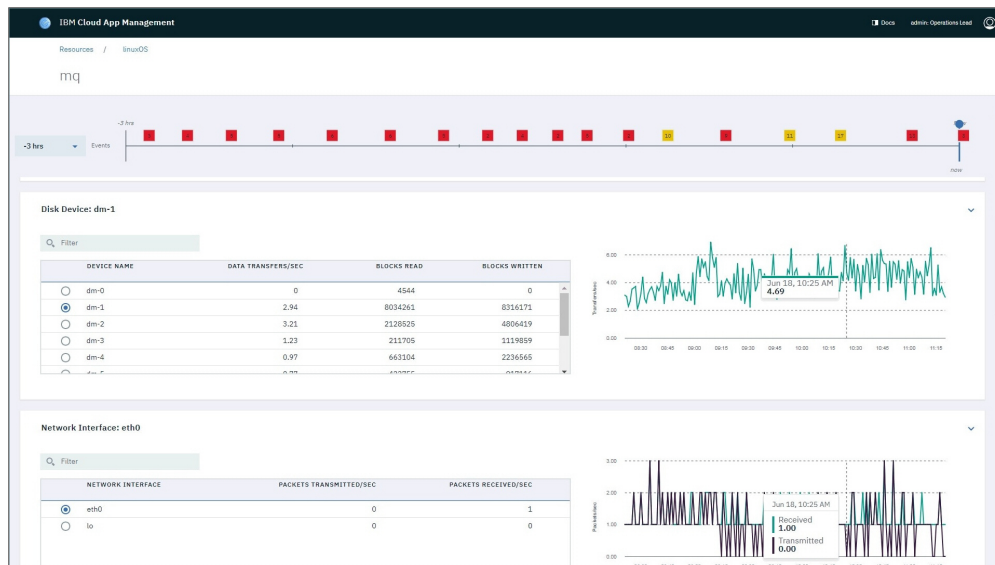
## User interface

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 shows 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

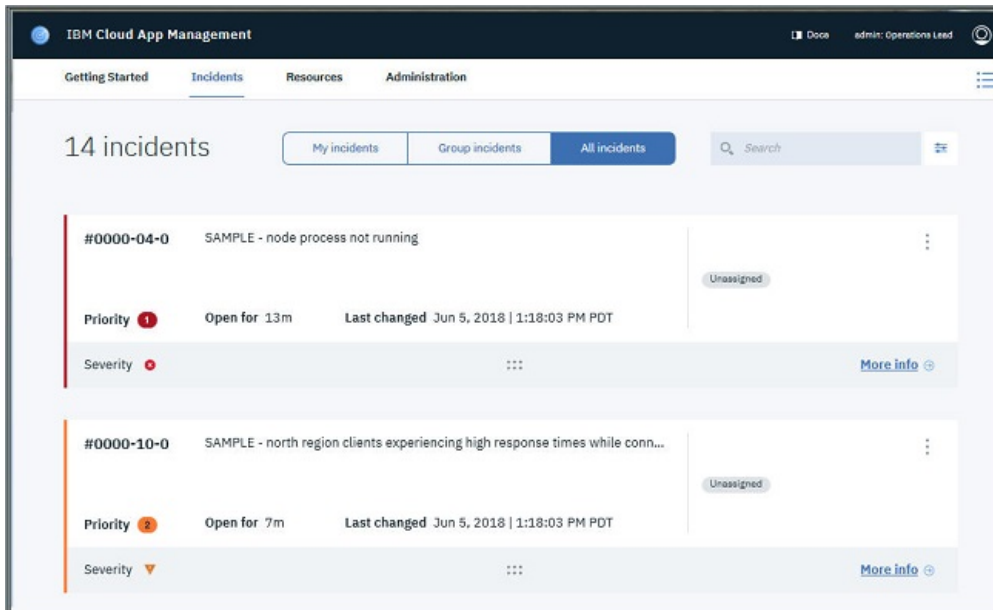
1. Select **Administration** from the menu bar.
2. Click the **Integrations** hexagon or its Information icon.



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:

#0001-16-0 SAMPLE - central region clients experiencing increased response times whil...  
 Priority 4 Open for 4m Last changed Jun 7, 2018 | 5:20:20 PM PDT  
 Owner Maureen  
 Group group1  
 Assigned

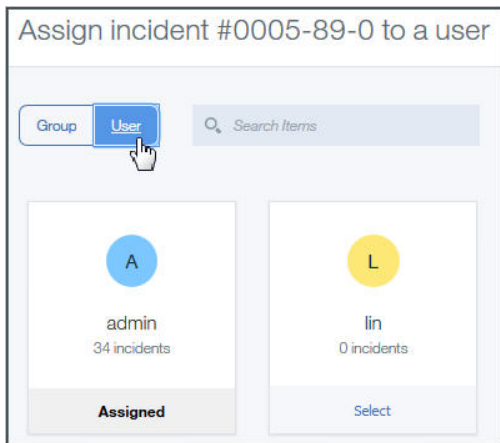
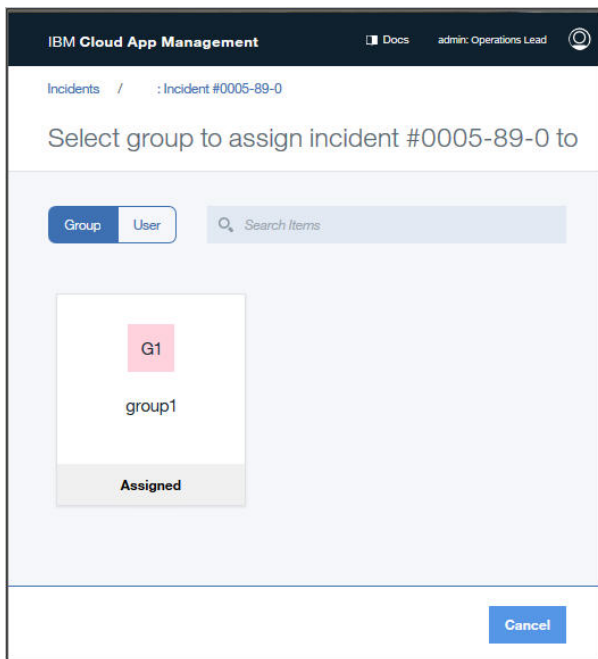
#0001-10-0 SAMPLE - Slow response time for application OnlineSales  
 Priority 4 Open for 10m Last changed Jun 7, 2018 | 5:20:19 PM PDT  
 Owner Maureen  
 Group group1  
 Assigned

### 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.

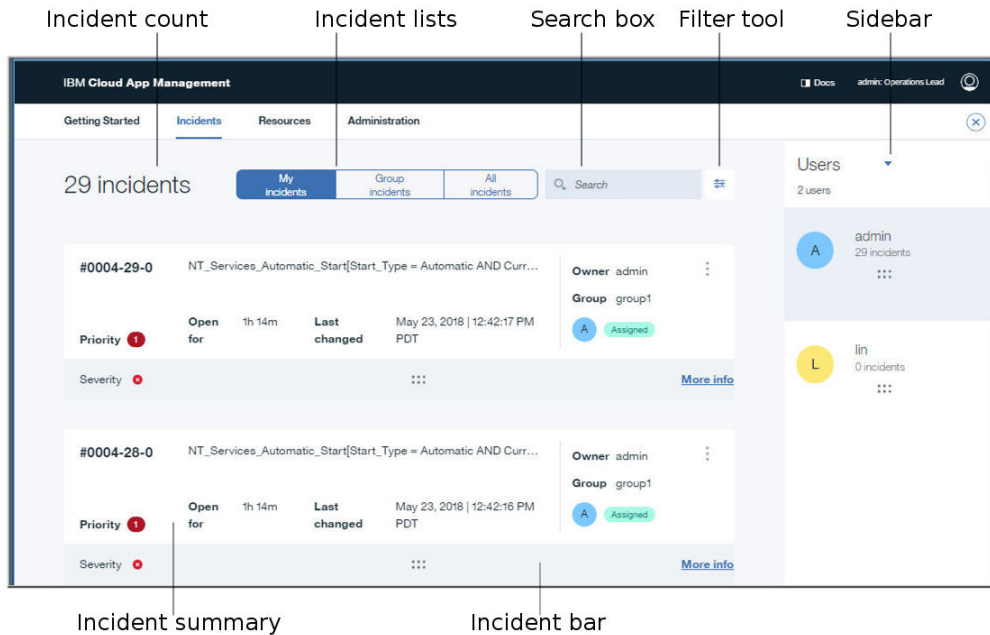
In this scenario, you'll assign one incident and work on another to find the root cause using the resource dashboards.

## 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.

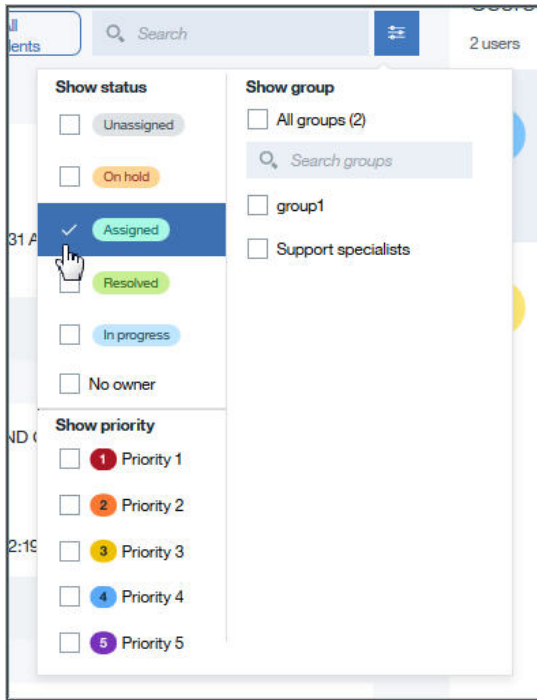
A summary view of each incident that is assigned to you is displayed.



## 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.



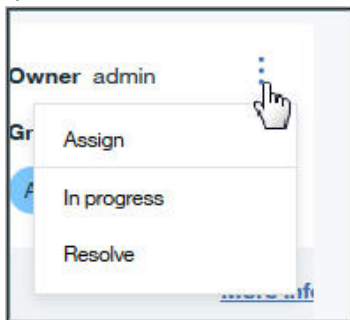
Two Priority 1 incidents are shown and both are assigned to another user. You read from the incident description that one incident is for high CPU usage and the other is for an inactive database.

### 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.



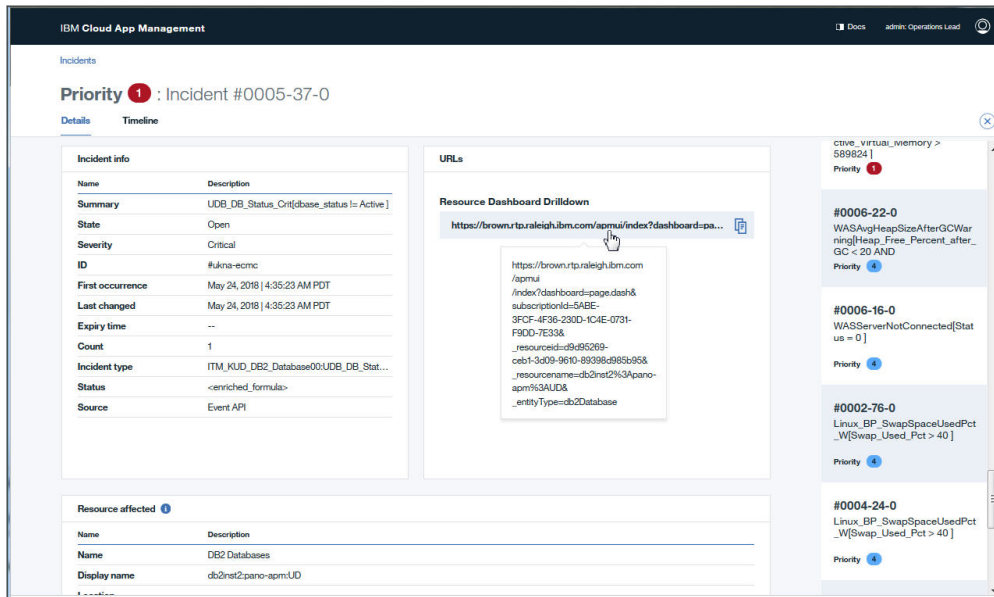
### Get more information about the incident

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

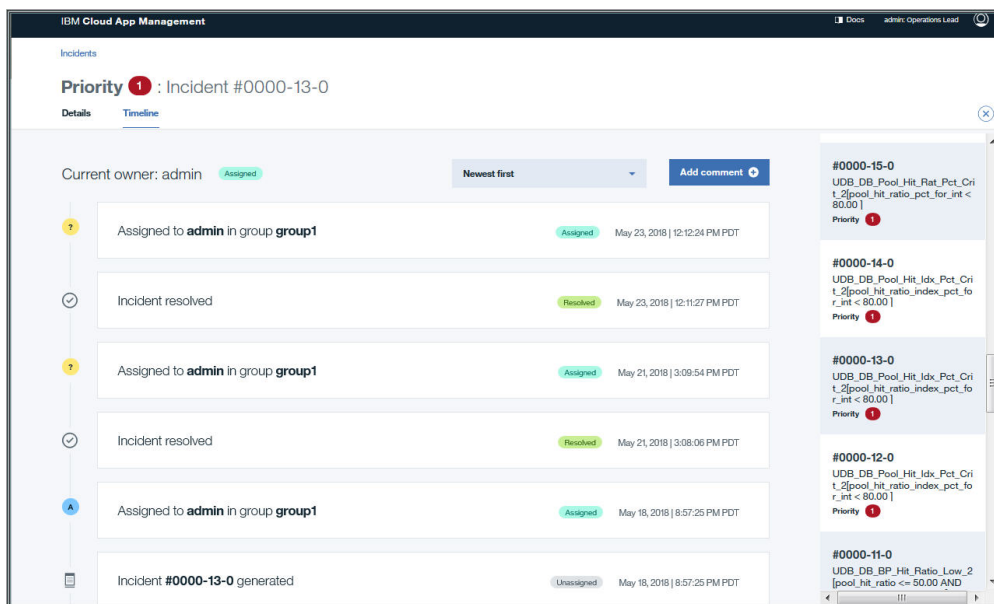
- The **Details** tab has information about the event and incident. The **First occurrence** and **Last changed** fields in the **Incident info** area tell when the threshold (or situation from your integrated Tivoli Monitoring agent) was first breached and the most recent change. Focus on this time 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.

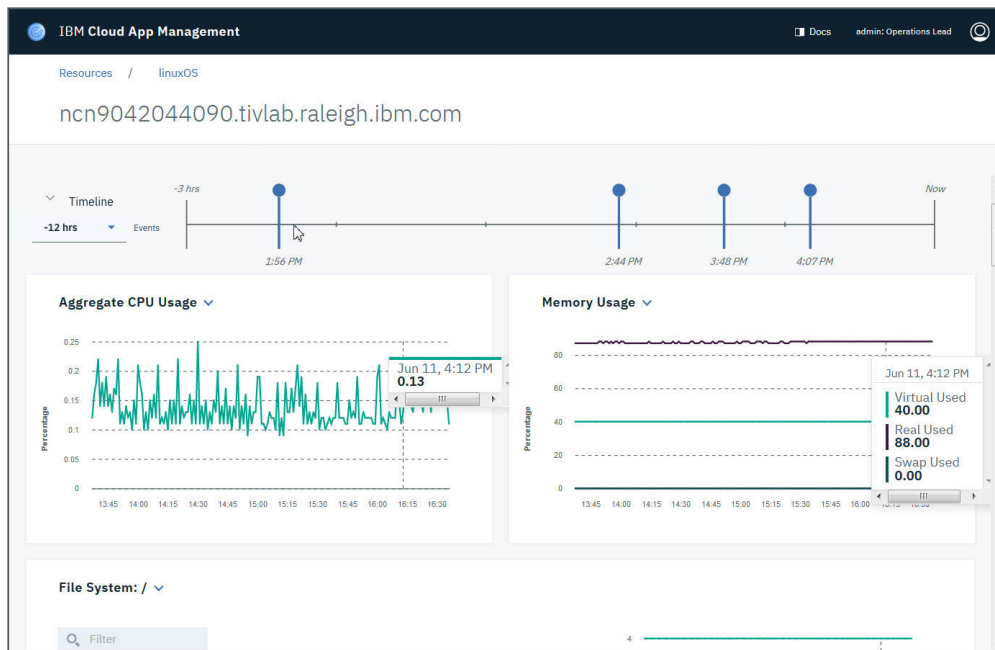


### Open the resource instance dashboard

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.

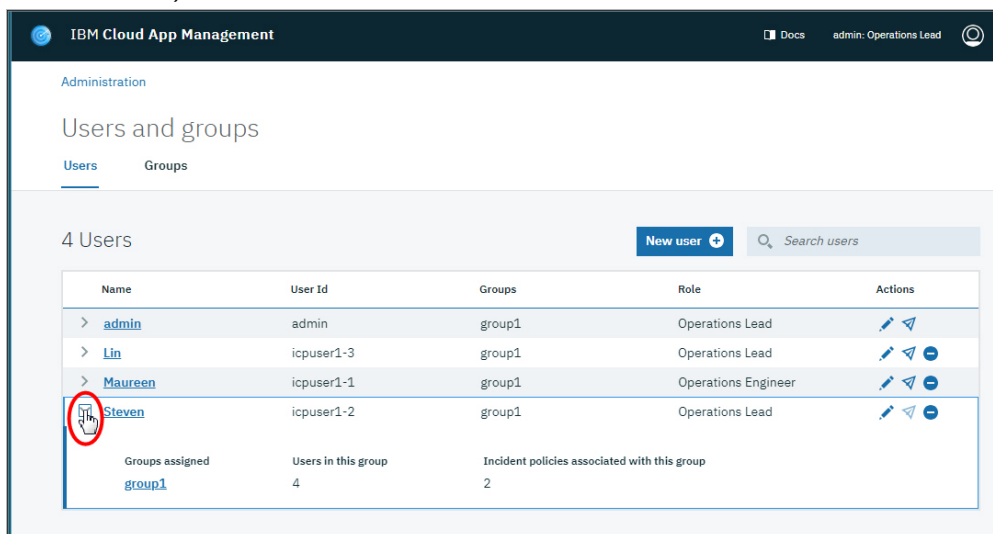
### Review user profiles

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 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 box labeled 'Filter threshold'. Below this is a table with the following 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

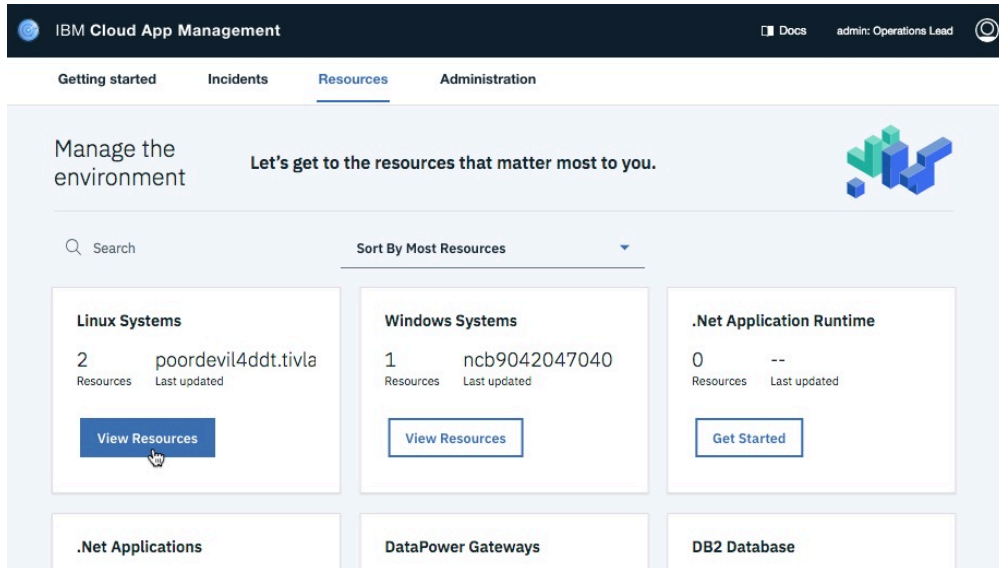
### Review the Resources dashboard

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.

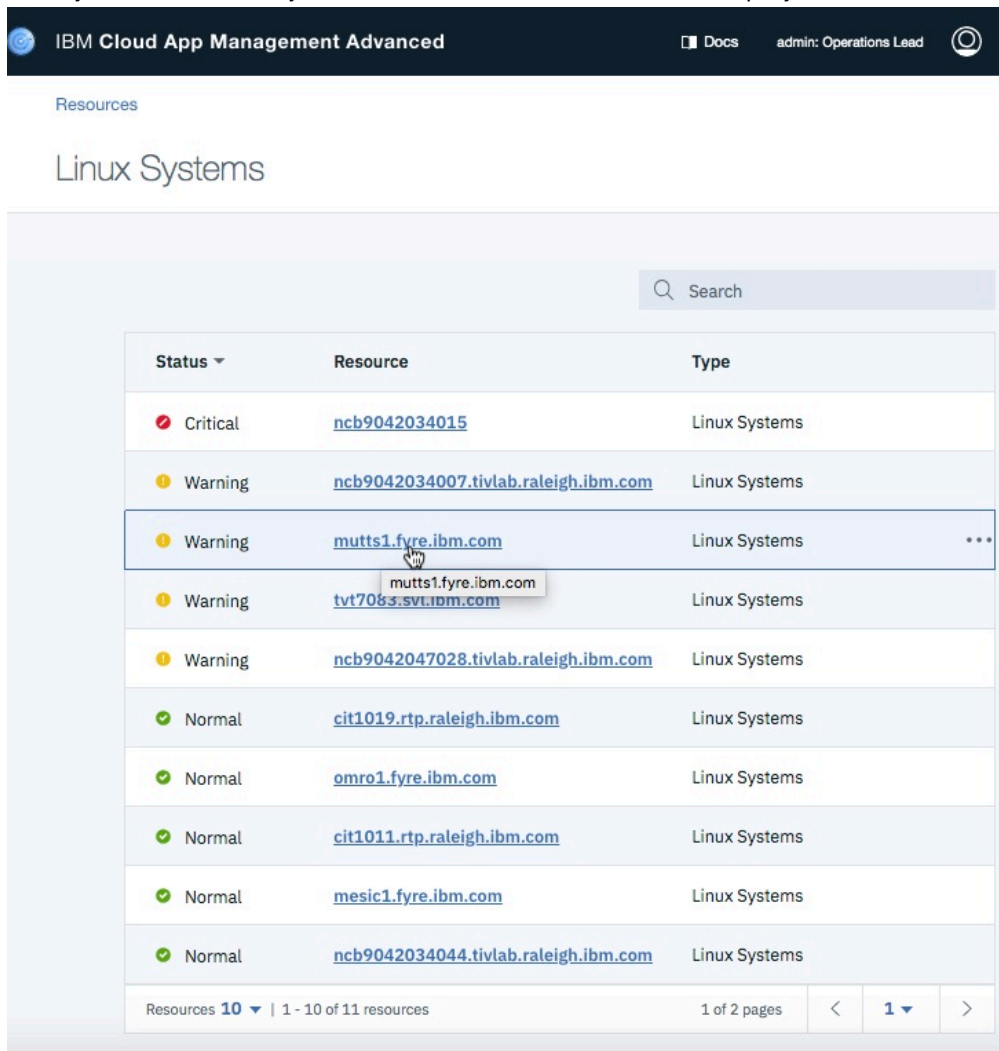
### All resources

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.

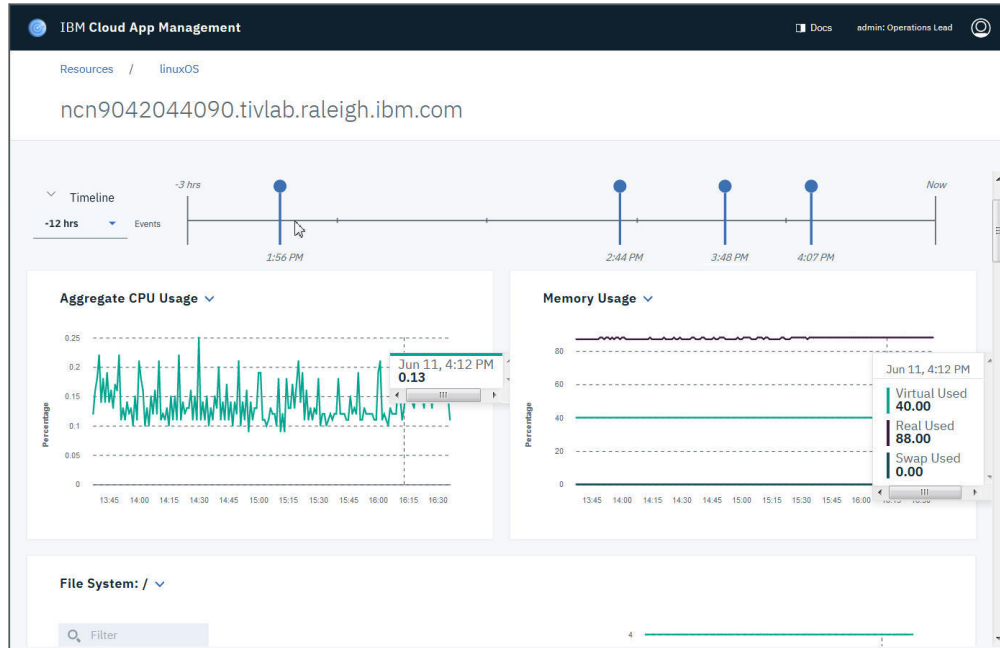




## Resource instance

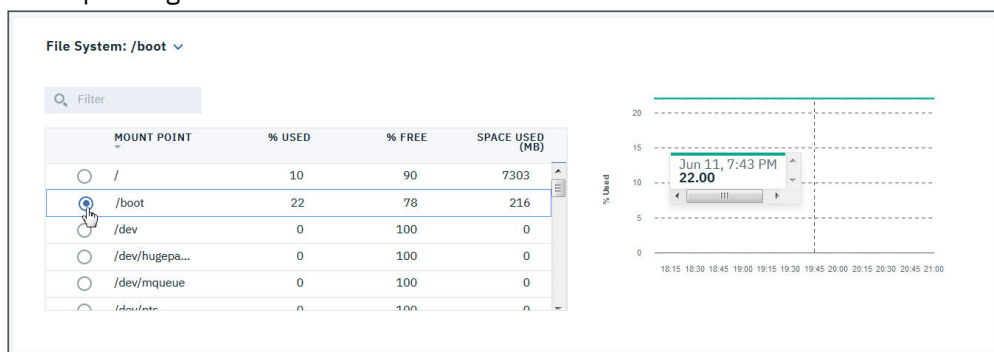
Select a link to open the dashboard for that instance. The instance dashboard is displayed with metrics from the past 3 hours by default. You can adjust the time span to show up to the past 32 days of saved samples from the data provider.

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.

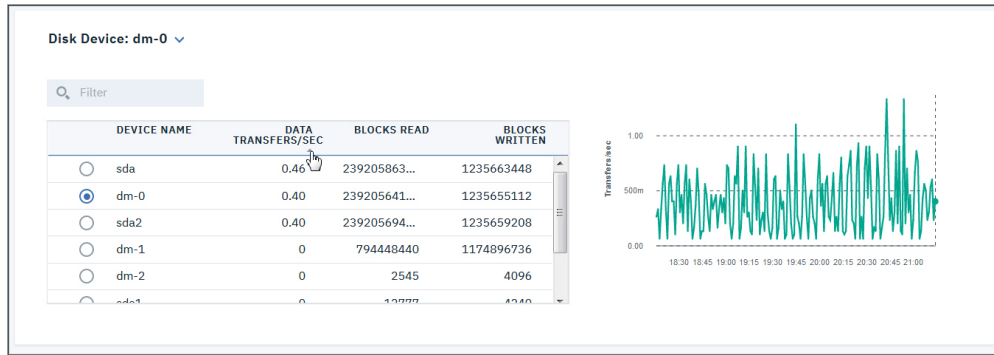


The Linux OS instance dashboard (also UNIX OS and Windows OS dashboards) plots metrics for the system characteristics. Dashboard sections with multiple metric views are synchronized:

- 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.



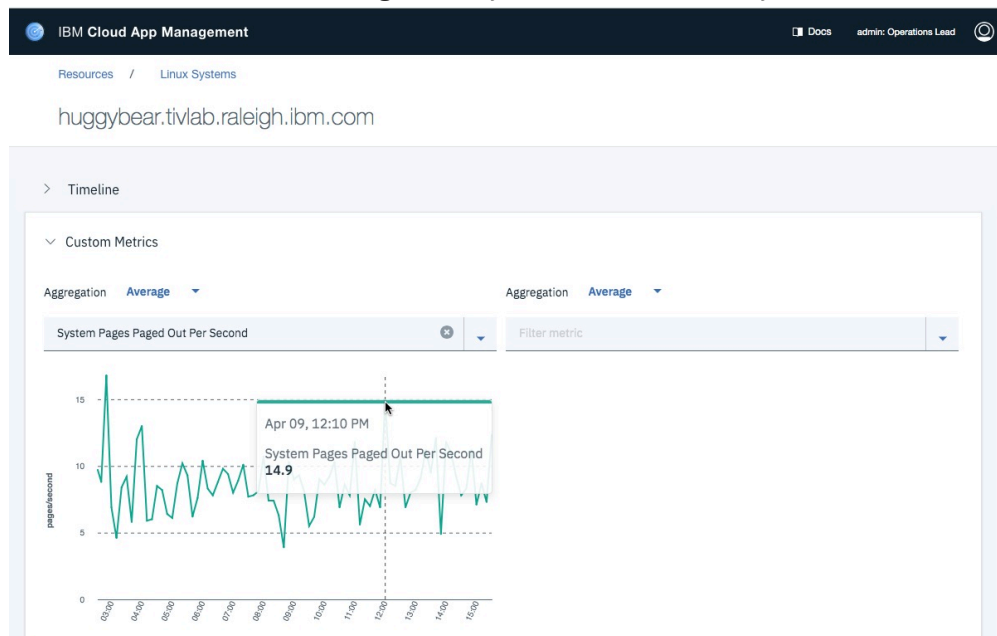
- Select the network interface type in the **Network Interface** table to see the number of packets transmitted and received per second plotted on the line chart.



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.

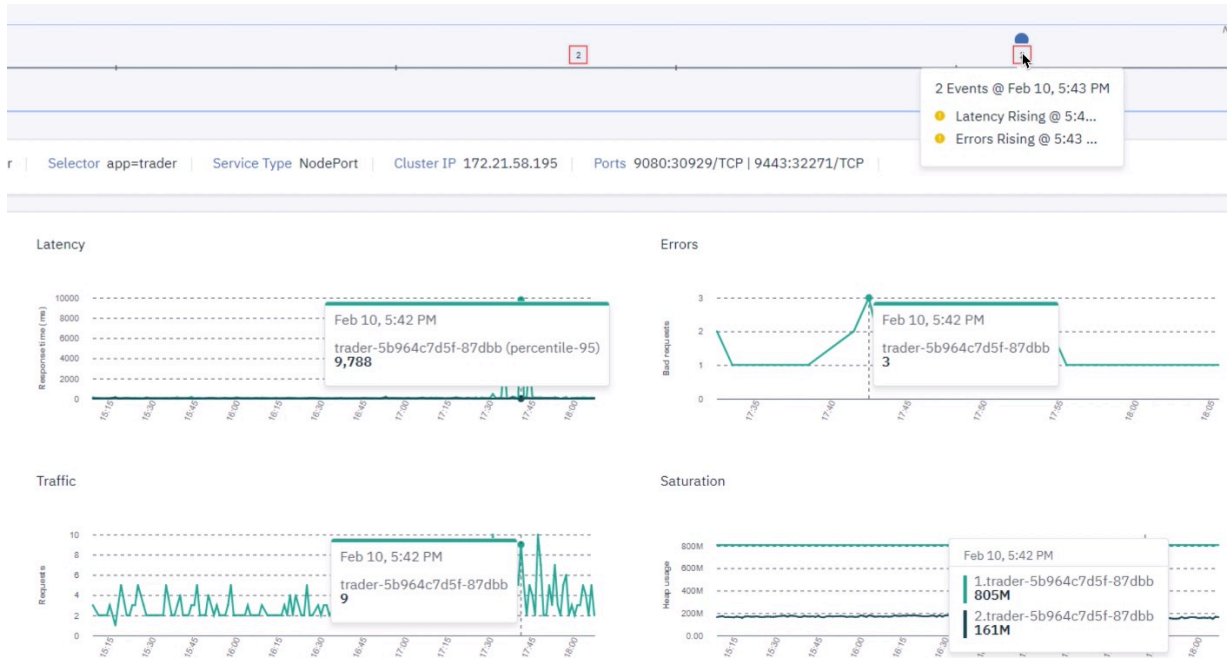
Typically, we would want to see if the problem is associated with a specific request type (or more) to determine where the root cause of the problem resides. We want to filter the signals by request to see which were impacted, and we notice that the latency increase is consistent across all request types;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

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

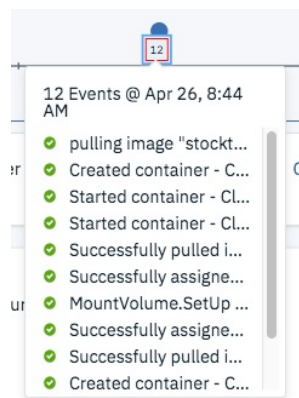
From the service 1-hop view, we see that Trader depends only on Portfolio. In this case, all requests go through Portfolio.



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 1

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

Administration

Runbooks

Library Execution Automations Connections

- 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.

```
[kubectl scale deployment portfolio --replicas=3]
```

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.

## Agents and data collectors

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.

Each agent and data collector 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 38](#). To find out the change history of each agent and data collector, see [“Change history” on page 37](#). For instructions on installing the agents and data collectors, see [Chapter 9, “Deploying ICAM Agents,” on page 131](#) and [Chapter 11, “Deploying ICAM Data Collectors,” on page 427](#).

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

The following agents and data collectors are supported in Cloud App Management:

- Cisco UCS agent
- DataPower® agent
- Db2 agent
- DataStage agent
- Hadoop agent
- HTTP Server agent
- IBM Integration Bus agent
- JBoss agent
- J2SE data collector
- Kubernetes data collector
- Liberty data collector
- Linux OS agent
- Linux KVM 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
- Node.js data collector
- Oracle Database agent
- PostgreSQL agent
- Python data collector
- SAP agent
- SAP HANA Database agent\*
- SAP NetWeaver Java Stack agent
- Skype for Business Server agent
- Synthetics PoP
- Tomcat agent
- Unified Agent

- UNIX OS agent
- VMware VI agent
- WebSphere Applications agent
- WebSphere Infrastructure Manager agent
- IBM MQ(formerly WebSphere MQ) agent
- Windows OS agent
- \* Extension pack agents

## Change history

Find out the information about versions and change history for each agent and data collector.

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

<i>Table 2. Agent and data collector change history</i>	
<b>Agents and data collectors</b>	<b>Links</b>
Cisco UCS 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>
Hadoop agent	<a href="#">Change history</a>
HTTP Server agent	<a href="#">Change history</a>
Cloud 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>
Liberty data collector	<a href="#">Change history</a>
Linux KVM agent	<a href="#">Change history</a>
Linux OS 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>
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>

<i>Table 2. Agent and data collector change history (continued)</i>	
<b>Agents and data collectors</b>	<b>Links</b>
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>
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>
Tomcat agent	<a href="#">Change history</a>
UNIX OS agent	<a href="#">Change history</a>
VMware VI agent	<a href="#">Change history</a>
WebSphere Applications agent	<a href="#">Change history</a>
WebSphere Infrastructure Manager agent	<a href="#">Change history</a>
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.

### 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 168.

### 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 11, “Deploying ICAM Data Collectors,”](#) on page 427.
- For information about the **Resource** dashboards, see [“Viewing your managed resources”](#) on page 611.



## 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 174.

## 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 178.

## 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 189.

## 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 197.

## 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 200.

## 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 209.

## 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 212.

## 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 465](#).

## 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 220](#).

## 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.

## 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 233](#)

## 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 231](#).

## 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 235.

### **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 246.

### **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 250.

### **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 252.

### **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 259.

### **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 261.

### **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 290.

## 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 295](#).

## 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 on the IBM Cloud Application Management UI.

For information about configuring the agent after installation, see [“Configuring NetApp Storage monitoring” on page 298](#).

## 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 450](#).

## 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 303](#).

## 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 339](#).

## Python monitoring

The Python data collector is a runtime data collector that helps you monitor Django based 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 470](#).

## 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 480](#).

## **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 345](#).

## **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 374](#).

## **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 377](#).

## **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 380](#).

## **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 385](#)

## **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 57](#).) 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 389.

### **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 397.

### **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 421.

### **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 421.

### **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 4. 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 \(www.ibm.com/able\)](http://www.ibm.com/able).

---

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





---

## Chapter 5. 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 6. 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 51](#)
- [“Step 2: Determine the storage type to use” on page 51](#)
- [“Step 3: Deploy IBM Cloud Private Enterprise V3.2.1” on page 52](#)
- [Step 4: Download and deploy the Cloud App Management server](#)
- [“Step 5: Access the Cloud App Management console” on page 52](#)
- [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 [<link>](#), 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 58](#) 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 to 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 7, “Deploying IBM Cloud Private,”](#) on page 69 section.

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

Complete the procedures in the [“Installing IBM Cloud App Management on IBM Cloud Private”](#) on page 99 section.

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

Complete the steps in the [“Starting the Cloud App Management UI”](#) on page 124 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 517 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 530 section.

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

Complete the procedures in the [Chapter 9, “Deploying ICAM Agents,”](#) on page 131 topics.

For configuring ICAM Data Collectors, complete the procedures in the [Chapter 11, “Deploying ICAM Data Collectors,”](#) on page 427 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 [“Downloading the Cloud App Management Passport Advantage Archive \(PPA\) file”](#) on page 100.
- If you plan to integrate with IBM Tivoli Monitoring agents, you must download the 6.3.0.7-TIV-ITM\_TEMA-IF0006 agent patch from IBM Fix Central, for more information, see [“Connecting IBM Tivoli Monitoring agents to Cloud App Management server”](#) on page 518.
- Download the agents installation images from IBM Passport Advantage. For more information, see [“Downloading agents and data collectors from Passport Advantage”](#) on page 132.

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 53](#).

## Part numbers

Review the part numbers to identify the components to download from the [IBM Passport Advantage](#) website for your IBM Cloud App Management V2019.3.0 installation.

### IBM Cloud App Management components

The part numbers and file names for the IBM Cloud App Management components are in [Table 3 on page 53](#).

eImage descriptions	IBM Cloud App Management (CJ618EN)
IBM Cloud App Management V2019.3.0 Server Install on AMD64	CC3FIEN app_mgmt_server_2019.3.0.tar.gz
IBM Cloud App Management V2019.3.0 Agents Install xLinux	CC3FJEN appMgtAgents_xlinux_2019.3.0.tar.gz
IBM Cloud App Management V2019.3.0 Agents Install zLinux	CC3FNEN appMgtAgents_zlinux_2019.3.0.tar.gz
IBM Cloud App Management V2019.3.0 Agents Install Windows	CC3FKEN appMgtAgents_win_2019.3.0.zip
IBM Cloud App Management V2019.3.0 Agents Install AIX	CC3FLEN appMgtAgents_aix_2019.3.0.tar.gz
IBM Cloud App Management V2019.3.0 Agents Install Solaris Sparc	CC3FPEN appMgtAgents_solaris_2019.3.0.tar.gz
IBM Cloud App Management V2019.3.0 Agents Install PLinux	CC3FQEN appMgtAgents_plinux_2019.3.0.tar.gz
IBM Cloud App Management V2019.3.0 Agents Install PLinuxLE	CC3FREN appMgtAgents_plinuxle_2019.3.0.tar.gz
IBM Cloud App Management V2019.3.0 Data Collectors Install	CC3FMEN appMgtDataCollectors_2019.3.0.tar.gz (Contains sub packages, described in <a href="#">“Data collectors sub-packages” on page 54</a> .)
Multicluster Event Management Klusterlet on PLinuxLE	CC3FSEN agent_ppa_2019.3.0_prod.tar.gz
Multicluster Event Management Server on PLinuxLE	CC3FTEN icam_ppa_2019.3.0_prod_lite.tar.gz
Multicluster Event Management Klusterlet on AMD64	CC3FUEN agent_ppa_2019.3.0_prod_amd64.tar.gz
Multicluster Event Management Server on AMD64	CC3FVEN icam_ppa_2019.3.0_prod_lite_amd64.tar.gz
IBM Cloud Unified Agent V2019.3.0	CC3FWEN unifiedAgent_2019.3.0.tar.gz

### Data collectors sub-packages

appMgtDataCollectors\_2019.3.0.tar.gz contains the following 3 sub-packages:

- app\_mgmt\_k8sdc.tar.gz (contains sub-packages)
- app\_mgmt\_runtime\_dc\_2019.3.0.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\_docker.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-IF0006 agent patch from IBM Fix Central, for more information, see [“Connecting IBM Tivoli Monitoring agents to Cloud App Management server”](#) on page 518.

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

eImage descriptions	IBM Cloud App Management V2019.3.0 Extension Pack (CJ619EN)
IBM Cloud App Management V2019.3.0 Extension Pack xLinux	CC3FXEN appMgtExt_xlinux_2019.3.0.tar
IBM Cloud App Management V2019.3.0 Extension Pack Windows	CC3FYEN appMgtExt_win_2019.3.0.zip
IBM Cloud App Management V2019.3.0 Extension Pack AIX	CC3FZEN appMgtExt_aix_2019.3.0.tar
IBM Cloud App Management V2019.3.0 Extension Pack PLinux	CC3G1EN appMgtExt_plinux_2019.3.0.tar
IBM Cloud App Management V2019.3.0 Extension Pack PLinuxLE	CC3G2EN appMgtExt_plinuxle_2019.3.0.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 54](#).

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>



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

eImage descriptions	IBM Cloud Private (CJ5NFEN)	Product documentation
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 55](#).

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 55](#) (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

Title on Passport Advantage	eAssembly number	Product documentation
IBM Cloud Application Performance Management, Base Private V8.1.4 for IBM Cloud App Management Multiplatform Multilingual eAssembly	CJ40EML	<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)*

<b>Title on Passport Advantage</b>	<b>eAssembly number</b>	<b>Product documentation</b>
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>
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 58.

### 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 Kubernetes data collector

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:

<b>Agents and data collectors</b>	<b>System requirements</b>
Cisco UCS agent	<a href="#">Cisco UCS 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>
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>
Linux OS agent	<a href="#">Linux OS agent</a>
Linux KVM agent	<a href="#">Linux KVM agent</a>
Microsoft .NET agent	<a href="#">Microsoft .NET 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>

<i>Table 8. Supported operating systems for agents (continued)</i>	
<b>Agents and data collectors</b>	<b>System requirements</b>
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>
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>
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>
Runtime data collectors	<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>
Synthetics PoP	<a href="#">Synthetics PoP</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 59 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 63 .

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 59 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: [Load projection spreadsheet](#) . The spreadsheet returns an estimate of total metrics per minute. For more information, see the [Using the IBM Cloud App Management Database Load Projections Spreadsheet](#) PDF .
2. Review [Table 9 on page 59](#) 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_amd64, --size0_ppc64le, --size 1_amd_64, --size1_ppc64le
```

For IBM Power installations, you can select **--size0\_ppc64le** or **--size1\_ppc64le**. This reduces the CPU allocation to reflect the generally lower CPU required by IBM Power. If for any reason, your IBM Power environment does require additional CPU, you can alternatively choose the **--size1\_amd64** or **--size1\_ppc64le** parameter. For detail of CPU requirement, see the [Cores and Memory row](#) in [Table 1](#) below.

For more information on how to specify size, see [step “8” on page 114](#) in the Cloud App Management server installation topic.

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

Table 9. Sizes (continued)

Category	Resource	Demo/POC/Trial		Entry		Standard	Enterprise
Total cores and Memory	AMD (default) CPUs	12	24	35	70	90	105
	IBM Power CPUs	6	12	18	35	45	53
	Memory (GB)	32	64	55	130	180	230
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

**Remember:**

The Total 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.

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 63.

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 100.

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 100 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:

<i>Table 10. Default ports used by agents</i>		
<b>Agents</b>	<b>Default ports</b>	<b>Configurable</b>
Cisco UCS agent	N/A	N/A
DataPower agent	5550 (for connecting to remote DataPower appliances)	Yes
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
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
Microsoft .NET agent	N/A	N/A

Table 10. Default ports used by agents (continued)

Agents	Default ports	Configurable
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
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
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
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
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



Agents	Default ports	Configurable
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 100](#).

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 64](#).

For Kafka, see [“Scaling up Kafka brokers” on page 66](#).

For CouchDB, see [“Scaling up CouchDB” on page 65](#).

For Datalayer, see [“Scaling up Datalayer” on page 68](#).

For ZooKeeper, see [“Scaling up ZooKeeper” on page 67](#).

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

- agentbootstrap
- agentmgmt
- alarmeventsrsc
- 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 58](#).

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 107](#).

## 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 \  
after -cassandra \ # Add the PV number immediately  
  --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 67, 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-zookeeperpv_number \ # Add the PV number immediately
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 7. 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 “[System requirements](#)” on page 57 topic. 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 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 Knowledge Center.

**Note:** After you install IBM Cloud Private Enterprise, you can add or remove management, worker, and proxy nodes from your cluster.
  - b. Prepare each node for installation. For more information, see the [Configuring your cluster](#) topic in the IBM Cloud Private 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 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 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 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 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 69 topic.

### About this task

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

```

- 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

```

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

```

docker info
systemctl status docker

```

- 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 Knowledge Center.
- 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 user name 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

```

- 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 Docker CLI](#) topic in the IBM Cloud Private 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 61](#).
- 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 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 Knowledge Center.
- Deploy the Cloud App Management server and agents. For more information, see the following topics: [“Installing IBM Cloud App Management on IBM Cloud Private” on page 99](#), [Chapter 9, “Deploying ICAM Agents,” on page 131](#), and [“Starting the Cloud App Management UI” on page 124](#).

## 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 8. Installing IBM Cloud App Management - the options

You can install IBM Cloud App Management in three different ways: with IBM Cloud Pak for Multicloud Management, on IBM Cloud Private, or on Red Hat OpenShift. Refer to the table in this topic for more details.

Installation option	Supported platforms and software required
<a href="#">“Installing IBM Cloud App Management with IBM Cloud Pak for Multicloud Management” on page 73</a>	<ul style="list-style-type: none"><li>• Linux<sup>®</sup> x86_64 with Red Hat OpenShift V3.11 and Cloud Pak Foundation V3.2.1</li><li>• Linux<sup>®</sup> on Power<sup>®</sup> with Red Hat OpenShift V3.11 and Cloud Pak Foundation V3.2.1</li></ul>
<a href="#">“Installing IBM Cloud App Management on Red Hat OpenShift” on page 91</a>	<ul style="list-style-type: none"><li>• Linux<sup>®</sup> x86_64 with Red Hat OpenShift V3.11 and Cloud Pak Foundation V3.2.1</li></ul>
<a href="#">“Installing IBM Cloud App Management on IBM Cloud Private” on page 99</a>	<ul style="list-style-type: none"><li>• Linux<sup>®</sup> x86_64 with IBM Cloud Private V3.11</li></ul>

### 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.

#### About this task

Installing IBM Cloud App Management with IBM Cloud Pak for Multicloud Management includes the following main steps:

1. [“Onboarding LDAP users” on page 73](#)
2. [“Deploying the Cloud App Management server with IBM Cloud Pak for Multicloud Management” on page 75](#)
3. [“Deploying the ICAM klusterlet for IBM Multicloud Manager ” on page 83](#)
4. [“Deploying agents and data collectors for IBM Multicloud Manager” on page 89.](#)
5. [“Uninstalling IBM Cloud App Management for IBM Multicloud Manager” on page 90.](#)

#### 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 you managed cluster to the team.

#### Before you begin

You must set up an LDAP connection in IBM Cloud Pak for Multicloud Management, select **Local Cluster>Identity & Access>Authentication**, and click **Create LDAP connection**. For more information, see [Configuring LDAP connection](#).

## 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 imports all the LDAP users in to IBM Cloud Pak for Multicloud Management.

Before you import users you can find 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 will use this ID in the next 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 74 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
```

Log in to IBM Cloud Pak for Multicloud Management.

4. 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://9.46.73.51:8443 -p user1-password -u user1ID
Authenticating...
OK
Targeted account testacc

User does not have any namespaces.
Enter a namespace name to be created:> my-namespace
Targeted namespace my-namespace

Configuring kubectl ...
Property "clusters.allowing-spider-icp-cluster" unset.
Property "users.allowing-spider-icp-cluster-user" unset.
Property "contexts.allowing-spider-icp-cluster-context" unset.
Cluster "allowing-spider-icp-cluster" set.
User "allowing-spider-icp-cluster-user" set.
Context "allowing-spider-icp-cluster-context" created.
Switched to context "allowing-spider-icp-cluster-context".
OK

Configuring helm: /Users/user1@ibm.com/.helm
OK
```

Create a team

5. Create a team for the ACOUNTID 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 74 into 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 into the team you created.

7. In IBM Cloud Pak for Multicloud Management select **Local Cluster>Identity & Access>Teams**. Select the team that you created in step “5” on page 75. Select the **Resources** tab, and click **Manage Resources**. Select your managed cluster and click **Save**. For more information, see [Add a Helm repository and namespace to a team](#).

## Deploying the Cloud App Management server with IBM Cloud Pak for Multicloud Management

These steps explain how to install and configure the Cloud App Management server with IBM Cloud Pak for Multicloud Management on the hub cluster.

### Before you begin

- You must install the Cloud Pak Foundation. For instructions, see the "Install IBM Multicloud Manager with Cloud Pak Foundation" section in the [Installing the IBM Cloud Pak for Multicloud Management](#) topic.
- To install IBM Cloud App Management with IBM Cloud Pak for Multicloud Management, you must import a target manager cluster into the IBM Multicloud Manager hub cluster. For more information, see [Importing a target managed cluster to the IBM Multicloud Manager hub cluster](#).
- You must onboard LDAP users. For more information, see [“Onboarding LDAP users” on page 73](#).

### Procedure

Complete the following steps as an IBM Cloud Private cluster administrator on your hub cluster:

1. Locate the Cloud App Management server installation image file `app_mgmt_server_2019.3.0.tar.gz` (part number: CC3FIEN) on [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 the [“Part numbers” on page 53](#).

**Note:** The Cloud App Management server with IBM Cloud Pak for Multicloud Management must be installed in the `kube-system` namespace.

2. As the IBM Cloud Private administrator, log in to the management console and select your namespace. Run the following command:

```
cloudctl login -a my_cluster_URL -n kube-system -skip-ssl-validation
```

Where `my_cluster_URL` is the IBM Cloud Private name that you defined for your cluster such as `https://cluster_address:443`. For any ‘masterIP’ reference, use the `cluster_address` value.

**Note:** Run the following command to get your cluster information before you start.

```
kubectl get configmap -n kube-public -o yaml
```

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)
```

#### 5. Create an install directory and download the IBM Passport Advantage file: `app_mgmt_server_2019.3.0.tar.gz` to the install directory. Change to the install directory:

```
mkdir -p install
cd install
```

#### 6. Extract the Helm charts from the Passport Advantage Archive (PPA) file into the `install_dir` directory. Extract the (PPA) file by running the following example commands:

```
cd install_dir
tar -xvf ./ppa_file charts
tar -xvf ./charts/ibm-cloud-appmgmt-prod-1.5.0.tgz
```

Where:

- `ppa_file` is the compressed Cloud App Management PPA installation image file: `app_mgmt_server_2019.3.0.tar.gz` file. The `charts` value is required to ensure the `tar` command extracts only the `charts` directory from the `ppa_file`. Otherwise, all the images are extracted, which might cause space issues.

#### 7. Load the PPA file into IBM Cloud Private Docker registry:

```
cloudctl catalog load-archive --archive ./ppa_file --registry $(oc registry info)/kube-system
```

#### 8. Change directory to the `ibm-cloud-appmgmt-prod` directory.

```
cd install_dir/ibm-cloud-appmgmt-prod
```

#### 9. Create the Cloud App Management ingress TLS and client secrets:

```
./ibm_cloud_pak/pak_extensions/lib/make-ca-cert-icam.sh my_ProxyHostName my_release_name kube-system
```

Where `my_release_name` is the name of the Cloud App Management Helm Chart, for example: `ibmcloudappmgmt`.

#### 10. Apply the OpenShift security context constraints into the namespace where you are installing Cloud App Management.

**Note:** This chart requires a `SecurityContextConstraints` to be bound to the target namespace before you install Cloud App Management. To meet this requirement, cluster and namespace preinstallation and postinstallation steps might need to be completed. The predefined `SecurityContextConstraints` name: `ibm-restricted-scc` is verified for this chart. If your target namespace is bound to this `SecurityContextConstraints` resource, you can start to install the chart. This chart also defines a custom `Constraints`, which can be used to control the permissions and capabilities that are needed to deploy this chart. You can enable this custom `SecurityContextConstraints` resource by using the following scripts.

- a) Run the following setup script as a cluster administrator. The postinstallation instructions are included in this script.

```
cd ibm_cloud_pak/pak_extensions/pre-install/clusterAdministration
./createSecurityClusterPrereqs.sh
```

- b) Run the following setup script as a team administrator. The namespace instructions are included in this script.

```
cd ibm_cloud_pak/pak_extensions/pre-install/namespaceAdministration
./createSecurityNamespacePrereqs.sh kube-system
```

c) Change directories back to the `install_dir/ibm-cloud-appmgmt-prod` directory. For example: `cd /install/ibm-cloud-appmgmt-prod`

11. Use the following command to create the ConfigMap for the cluster certificate authority (CA): `icam-cluster-ca-cert`:

```
cat <<EOF | kubectl apply -f -
kind: ConfigMap
apiVersion: v1
metadata:
  name: icam-cluster-ca-cert
  namespace: kube-system
data:
  ca_cert.crt: |
$(kubectl -n kube-public get secret ibmcloud-cluster-ca-cert -o jsonpath='{.data.ca\.crt}'
| base64 -d | sed 's/^/ /')
EOF
```

12. You must create the Cassandra auth secret with the following requirements:

- The Cassandra auth secret must have two keys: `username` and `password`. You must use only alphanumeric characters for the password.
- It must be created in the same namespace (`kube-system`) as the Cloud App Management server.

```
kubectl create secret generic $RELEASENAME-cassandra-auth-secret -n kube-system --from-literal=username=$CASSANDRA_USER --from-literal=password=$CASSANDRA_PASS
```

Replace the `$RELEASENAME`, `$CASSANDRA_USER`, and `$CASSANDRA_PASS` variables with values, where `$RELEASENAME` is the name that you are using for the Cloud App Management release; for example: `ibmcloudappmgmt`, and `$CASSANDRA_USER` and `$CASSANDRA_PASS` are the user name and password that you enter for the Cassandra auth secret.

```
kubectl create secret generic $RELEASENAME-cassandra-auth-secret -n kube-system --from-literal=username=$CASSANDRA_USER --from-literal=password=$CASSANDRA_PASS
```

The following example sets the user name and password both to Cassandra:

```
kubectl create secret generic ibmcloudappmgmt-cassandra-auth-secret -n kube-system --from-literal=username=cassandra --from-literal=password=cassandra
```

13. Remove port 80 from the `nginx-ingress` service:

```
kubectl edit services -n kube-system nginx-ingress
```

Remove the lines that have a line through them below.

```
apiVersion: v1
kind: Service
metadata:
  annotations:
    kubernetes.io/last-applied-configuration: |
      {"apiVersion":"v1","kind":"Service","metadata":{"annotations":{},"name":"nginx-ingress","namespace":"kube-system"},"spec":{"ports":[{"name":"http","port":80,"protocol":"TCP","targetPort":80}, {"name":"https","port":443,"protocol":"TCP","targetPort":443}],"selector":{"app":"nginx-ingress-controller"},"type":"ClusterIP"}}
    creationTimestamp: 2019-05-17T18:44:00Z
  name: nginx-ingress
  namespace: kube-system
  resourceVersion: "54038"
  selfLink: /api/v1/namespaces/kube-system/services/nginx-ingress
  uid: bc20e4de-78d3-11e9-9373-00163e01c818
spec:
  clusterIP: 172.30.79.16
  ports:
  - name: http
  - port: 80
  - protocol: TCP
  - targetPort: 80
  - name: https
  - port: 443
```

```
protocol: TCP
targetPort: 443
selector:
  app: nginx-ingress-controller
sessionAffinity: None
type: ClusterIP
status:
  loadBalancer: {}
```

14. Create local directories on each of the IBM Cloud Private worker nodes. Record the IP address of each worker node where the directories are created. You need these values when you are selecting your storage class and preparing persistent volumes in the next step in this procedure. Review the optimization performance topics. For more information, see [“Optimizing performance” on page 107](#).

**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.

- /k8s/data/cassandra - Cassandra persistent storage

```
mkdir -p /k8s/data/cassandra
```

- /k8s/data/zookeeper - ZooKeeper persistent storage

```
mkdir -p /k8s/data/zookeeper
```

- /k8s/data/kafka - Kafka persistent storage

```
mkdir -p /k8s/data/kafka
```

- /k8s/data/couchdb - CouchDB persistent storage

```
mkdir -p /k8s/data/couchdb
```

- /k8s/data/datalayer - Datalayer persistent storage

```
mkdir -p /k8s/data/datalayer
```

- /k8s/data/elasticsearch - Elasticsearch persistent storage

```
mkdir -p /k8s/data/elasticsearch
```

15. Select your storage class and prepare the persistent volumes for each of the following Cloud App Management services by running the `prepare-pv.sh` script. Local storage is recommended. The services are:

- Cassandra
- Kafka
- ZooKeeper
- CouchDB
- Data layer
- Elasticsearch

**Note:** For additional information, see the "Storage" section in the **Overview** tab by selecting the IBM-cloud-appmgmt-prod tile from the catalog. For information about how to configure persistent storage, see the following topics: [“Planning hardware and sizing” on page 58](#), [Understanding Kubernetes storage](#), [Planning a storage solution](#), [Planning persistent storage](#), and [Storage options in IBM Cloud Private](#)

**Note:** Before you run the `prepare-pv.sh` script, determine the addresses of the worker nodes where the persistent storage directories are located for each service by running the following command:

```
kubectl get nodes
```



The addresses can be the IP or hostname depending on what IBM Cloud Private is using.

- For local storage, specify the **--local** parameter. Local storage is recommended.
- For vSphere, specify the **--vSphere** parameter.
- Choose from predefined **--size0\_amd64**, **--size0\_ppc64le**, **--size1\_amd64**, or **--size1\_ppc64le** storage values. For more information about sizes, see [“Planning hardware and sizing”](#) on page 58.

```
ibm_cloud_pak/pak_extensions/prepare-pv.sh
--size0_amd64      #Install as size0 on amd64 (minimum resource requirements) - if
omitted, specify each size in the size parameters below
--size0_ppc64le   #Install as size0 on ppc64le (minimum resource requirements) - if
omitted, specify each size in the size parameters below
--size1_amd64     #Install as size1 on amd64 (standard resource requirements) - For
HA use --size1_amd64 - if omitted, specify each size in the size parameters below
--size1_ppc64le   #Install as size1 on ppc64le (standard resource requirements) -
For HA use --size1_ppc64le - if omitted, specify each size in the size parameters below
--releasename     #installation name, such as ibmcloudappmgmt, as defined in your
cluster image policy.

#Required flags for local storage
--local           #Use local persistent volume storage
#For high availability, list the nodes in quotes, separated by spaces
#If IBM Cloud
Private uses the IP address instead of the hostname, the IP address is needed here.
--CassandraNodes #IP or hostname of the node/s with the persistent storage local
directory for the Cassandra service.
--ZookeeperNodes #IP or hostname of the node/s with the persistent storage local
directory for the Zookeeper service.
--KafkaNodes     #IP or hostname of the node/s with the persistent storage local
directory for the Kafka service.
--CouchDBNodes  #IP or hostname of the node/s with the persistent storage local
directory for the CouchDB service.
--DatalayerNodes #IP or hostname of the node/s with the persistent storage local
directory for the Datalayer service.
--ElasticsearchNodes #IP or hostname of the node/s with the persistent storage local
directory for the Elasticsearch service.

#Optional storage directory paths for local storage
--CassandraDir   #the local system directory for Cassandra (default is /k8s/data/
cassandra)
--CassandraBackupDir #the local system backup directory for Cassandra (default is /k8s/
data/cassandra_backup)
--KafkaDir       #the local system directory for Kafka (default is /k8s/data/kafka)
--ZookeeperDir   #the local system directory for Zookeeper (default is /k8s/data/
zookeeper)
--CouchDBDir     #the local system directory for CouchDB (default is /k8s/data/
couchdb)
--DatalayerDir   #the local system directory for Datalayer (default is /k8s/data/
datalayer)
--ElasticsearchDir #the local system directory for Elasticsearch (default is /k8s/
data/elasticsearch)

#Optional storage class name flags for local storage:
--CassandraClass #the storage class name for Cassandra (default is <release_name>-
local-storage-cassandra)
--CassandraBackupClass ##the storage class name for Cassandra backups (default is
<release_name>-local-storage-cassandrabackup)
--KafkaClass     #the storage class name for Kafka (default is <release_name>-
local-storage-kafka)
--ZookeeperClass #the storage class name for Zookeeper (default is <release_name>-
local-storage-zookeeper)
--CouchDBClass  #the storage class name for CouchDB (default is <release_name>-
local-storage-couchdb)
--DatalayerClass #the storage class name for Datalayer (default is <release_name>-
local-storage-datapayer)
--ElasticsearchClass #the 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:
#Persistent volumes are measured in bytes.
#When you specify the size, you must use only the number and prefix, for example, 50
Gi. If you specify bytes, such as 50 GiB, the installation fails.
--CassandraSize  #the size of persistent volume for Cassandra (default size0 is 50
```

```

Gi)
--CassandraBackupSize #the size of persistent volume for Cassandra backups (default
size0 is 50 Gi)
--KafkaSize           #the size of persistent volume for Kafka (default size0 is 10 Gi)
--ZookeeperSize      #the size of persistent volume for Zookeeper (default size0 is 1
Gi)
--CouchDBSize        #the size of persistent volume for CouchDB (default size0 is 1 Gi)
--DatalayerSize      #the size of persistent volume for Datalayer (default size0 is 1
Gi)
--ElasticsearchSize  #the size of persistent volume for Elasticsearch (default size0
is 1 Gi)

```

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"

```

All directories for a statefulset type defined by `prepare-pv.sh` are in the same directory, for example: `/k8s/data/cassandra/`. If you want to use different directories, after you run `prepare-pv.sh` you can customize the `.yaml` file. Complete any customization to the `yaml` file before you run the `kubectl create` command. By default the `yaml` files are generated in: `ibm-cloud-appmgmt-prod/ibm_cloud_pak/yaml/`

16. Configure the Elasticsearch `vm.max_map_count` parameter:

- a) For Elasticsearch, you must set a kernel parameter to run normally for all worker nodes where it runs. These nodes are identified when you configure the persistent storage. You need to set `vm.max_map_coun` to a minimum value of 1048575. Set the parameter with `sysctl` to ensure that the change takes effect immediately. For example:

```
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
```

17. Log in to the IBM Cloud Private management console of your target cluster.

18. Click **Catalog**.

19. Search for and select the **ibm-cloud-appmgmt-prod** Helm Chart.

20. Click **Configure**.

21. Expand **All parameters**, configure the following parameters.

#### Helm release name

Enter a Helm release name for the deployment.

**Target namespace**

From the drop-down menu, select the **kube-system** namespace.

**Target Cluster**

The cluster that the ibm-cloud-appmgmt-prod Helm chart is being installed on.

**License Accepted**

Select the checkbox to accept the license.

22. Configure the following required parameters. Select **All parameters**.

**Create CRD, Resource monitoring services enabled, and Event Monitoring services enabled parameters**

The **Create CRD** (create the custom resource definition) check box is enabled by default, do not change it. The **Resource monitoring services enabled** and the **Event Monitoring services enabled** are also enabled by default. When these check boxes are enabled, it means IBM Cloud App Management is installed in full monitoring mode.

**Ingress Domain**

The fully qualified domain name or the IP address of your IBM Cloud Private management console.

**Ingress Port**

This value must match the IBM Multicloud Manager ingress port. The default port: 443 is usually used.

**Ingress Client Secret Name**

This field must be left empty.

**Ingress TLS Secret Name**

This field must be left empty.

**Proxy Ingress Domain**

The fully qualified domain name of your IBM Cloud Private Proxy.

**Proxy Ingress Client Secret Name**

The client secret name that was created in [step 9](#). It includes client authorization data for Proxy ingress. For example: *release\_name-ingress-client* might be *ibmcloudappmgmt-ingress-client*.

**Proxy Ingress TLS Secret Name**

The TLS secret name that was created in [step 9](#). It includes the HTTPS TLS authorization data. For example: *releasename-ingress-tls* might be *ibmcloudappmgmt-ingress-tls*.

**ICP Master Certificate Authority**

Use the name that you used in [Step 11](#). For example: *icam-cluster-ca-cert*. If you did not provide your own certificate, then you can leave this field empty.

```
kubectl create configmap master-ca --from-file=./ca.pem
```

**Cluster Master IP**

The master IP address of the IBM Cloud Private console.

**Cluster Master Port**

The port that used to communicate with the master IP address of the IBM Cloud Private console. The default port: 443 is usually used because it is also used with ingress domain IP address value in the IBM Multicloud Manager deployment.

**Cluster Proxy IP**

The IP address of your IBM Cloud Private Proxy.

**Product Deployment Size**

Enter the deployment size based on the resource requests and limits for the product. Remember, the Linux<sup>®</sup> on Power<sup>®</sup> CPU requests and limits are only half that of the Linux<sup>®</sup> x86\_64 CPU requests and limits. **test** is size 0 and **production** is size 1.

**Image Prefix**

The prefix for the docker images. It applies after the image repository value and before the image names. This field must be left empty.

**Image Repository**

The image repository must be set to `docker-registry.default.svc:5000/kube-system`. Remove the trailing slash from the value.

**Cassandra Replicas**

Configure the Cassandra Replicas.

**Image Pull Secret Name**

The name of the image pull secret that includes the credentials for accessing a private docker registry. Configuring this field is optional. You can leave it empty.

**Image pull policy**

The default pull policy is **IfNotPresent**, which causes the kubelet to skip pulling an image when it already exists. It is an optional field. You can leave the default value as it is.

**Cassandra Replicas**

Ensure these are consistent with the configuration in [step 15](#).

**Persistent Storage**

•

If you follow the instructions in [step 15](#) to create the required Persistent Volumes, review the `.yaml` files to ensure that they are consistent with the settings here in these fields. Leave the storage class default if you did not change it from the output of `pv-create.sh` in [step 15](#), otherwise make it consistent to the storage class specified for each Persistent Volume.

23. In the **CEM Configuration** section, configure the following parameters:

**Product Name**

IBM Cloud Pak for Multicloud Management

**ICP Cluster administrator user**

Cluster administrator user name. The default name is `admin`, 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.

24. Configure the parameters for **IBM Redis**.

25. 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.

**What to do next**

Complete the following postinstallation steps to complete the OIDC registration:

1. Click **View Helm Release** on the **Installation started** window. If this window fails to display, from the **IBM Cloud Pak for Multicloud Management** Menu, select **Helm Releases**, and then select the release name that you gave for IBM Cloud App Management.
2. Complete **step 3 (OIDC registration)** of the Notes section that is displayed on-screen. Run the **kubect1** command that is displayed in step 3 in the Notes section.
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 Multicloud Manager environment. For more information, see [“Deploying the ICAM klusterlet for IBM Multicloud Manager”](#) on page 83 to deploy the ICAM klusterlet using its Helm Chart or [“Deploying the ICAM klusterlet for IBM Multicloud Manager without helm”](#) on page 85 for a non-Helm installation.

## Deploying the ICAM klusterlet for IBM Multicloud Manager

By installing the ICAM klusterlet on any managed clusters, you can use this ICAM klusterlet to configure Prometheus and Kubernetes events.

### 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, see [“Deploying the Cloud App Management server with IBM Cloud Pak for Multicloud Management”](#) on page 75.
- The ICAM klusterlet that you are installing in this procedure must be installed in the same namespace as IBM Multicloud Manager.

### About this task

Install the IBM Cloud CLI from the IBM Cloud Private 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

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 Multicluster Event Management Klusterlet for PlinuxLE (part number: CC3FSEN): `agent_ppa_2019.3.0_prod.tar.gz`
- For Linux® x86\_64, choose Multicluster Event Management Klusterlet on AMD64 (part number: CC3FUEN): `agent_ppa_2019.3.0_prod_amd64.tar.gz`

2. Extract the PPA installation image file:

```
tar -xvf ppa_file
```

Where *ppa\_file* is the ICAM klusterlet PPA image file name.

3. Load the ICAM klusterlet PPA installation image file into the IBM Cloud Private Docker registry:

- a) Log in to the Docker registry. If you don't know what your Docker registry is, run the `oc registry info` command first to find out.

```
docker login docker-registry-url -u $(oc whoami) -p $(oc whoami -t)
```

- b) **Important:** After you download the installation package, load the klusterlet package CC3FSEN or CC3FUEN (depending on your platform) to the `multicluster-endpoint` namespace.

Load the ICAM klusterlet PPA image file into the IBM Cloud Private docker registry:

```
cloudctl catalog load-archive --archive install_dir/ppa_file --registry docker-registry-url/my_namespace
```

Where *ppa\_file* is the ICAM klusterlet PPA image file name, *docker-registry-url* is the docker registry (If you don't know what your Docker registry is, run the `oc registry info` command first to find out), and *my\_namespace* is `multicluster-endpoint`.

4. You must configure the namespace that IBM Multicloud Manager created for the IBM Cloud Private cluster. Run the following command on the IBM Multicloud Manager hub cluster to get a list of namespaces on the system:

```
kubect1 get clusters --all-namespaces
```

5. Locate your managed cluster namespace and cluster name from the list that is provided in [step 4](#) and enter it into the global **ICP cluster namespace** and global **Cluster Name** attributes.

- The Helm release name of the IBM Multicloud Manager klusterlet on the IBM Cloud Private system must be set in the **MCM Fullname Override** field. In the example below, 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-2f6ml 1/1 Running 0 2m
kube-system example-ibm-mcm-klusterlet-weave-scope-48mmn 1/1 Running 0 2m
kube-system example-ibm-mcm-klusterlet-weave-scope-app-5dds5r 1/1 Running 0 2m
kube-system example-ibm-mcm-klusterlet-weave-scope-g8gn4 1/1 Running 0 2m
```

You can issue the following command to obtain the **fullnameOverride**:

```
kubectl get pods | grep klusterlet
```

If your IBM Multicloud Manager was installed via the IBM Cloud Private CLI (cloudctl) instead of the Helm Chart, run the following command on your IBM Multicloud Manager klusterlet to obtain the **fullnameOverride**.

```
kubectl get secrets -n multicluster-endpoint |grep hub-kubeconfig
endpoint-connmgr-hub-kubeconfig Opaque 1 8d
The fullnameOverride value is: endpoint-connmgr
```

- Log in to the IBM Cloud Private management console of your target cluster.
- Click **Catalog**.
- Search for and select the **icam-clouddc-klusterlet** Helm Chart.
- Click **Configure**.
- Configure the following **Configuration** parameters:
  - Helm release name**  
Enter a string value in Helm release name field.
  - Cluster name**  
The name that this cluster will be identified as on the hub cluster.
  - Target Cluster**  
The cluster on which icam-clouddc-klusterlet is being installed.
  - Target namespace**  
From the drop-down menu, select the namespace where the IBM Multicloud Manager klusterlet was installed.
  - MCM Fullname Override**  
The helm release name for the IBM Multicloud Manager klusterlet on the IBM Cloud Private system.
  - Images Repository**  
The image repository must be set to the cluster name and namespace that the PPA was loaded to.
- Click **Install** to deploy the icam-clouddc-klusterlet Helm chart.

## What to do next

- Update the IBM Multicloud Manager team for Cloud App Management. The namespace corresponding to any new managed cluster must be added as a resource to the Cloud App Management team. On the IBM Cloud Private console, click **Manage > Identity and Access > Teams**. For more information, see [Adding the Helm repository and namespace to a team](#). Click the **Event Management** menu item to update the ICAM klusterlet with the correct webhook for Cloud App Management events.
- 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 Multicloud Manager”](#) on page 89.

## Deploying the ICAM klusterlet for IBM Multicloud Manager 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

1. The Cloud App Management server with IBM Cloud Pak for Multicloud Management must be installed on the hub cluster. For more information, see [“Deploying the Cloud App Management server with IBM Cloud Pak for Multicloud Management”](#) on page 75.

### 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: CC3FSEN): `agent_ppa_2019.3.0_prod.tar.gz`
- For Linux® x86\_64, choose Multicloud Event Management Klusterlet on AMD64 (part number: CC3FUEN): `agent_ppa_2019.3.0_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 85.

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 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_201909180825.tar.gz
images/k8-monitor_APM_201909192049.tar.gz
images/k8sdc-operator_APM_201909210313.tar.gz
images/reloader_201906210402-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_201909192049.tar.gz
docker tag k8-monitor:APM_201909192049 my_cluster_ca_domain:my_docker_registry_port/
my_namespace/k8-monitor:APM_201909192049
docker push my_cluster_ca_domain:my_docker_registry_port/my_namespace/k8-
monitor:APM_201909192049
```

Where

*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

5. Create Docker **imagePullSecrets**:

```
kubectl config set-context my_cluster_ca_domain-context --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. Deploy resources for the ICAM klusterlet:

### a) Create the "custom resource definition":

```
kubectl create -f k8monitor_crd.yaml
```

A template of the `k8monitor_crd.yaml` file is shown here:

```
apiVersion: apiextensions.k8s.io/v1beta1
kind: CustomResourceDefinition
metadata:
  name: k8sdcs.ibmcloudappmgmt.com
spec:
  group: ibmcloudappmgmt.com
  names:
    kind: K8sDC
    listKind: K8sDCList
    plural: k8sdcs
    singular: k8sdc
  scope: Namespaced
  subresources:
    status: {}
  version: v1alpha1
  versions:
  - name: v1alpha1
    served: true
    storage: true
```

### b) Apply any necessary value changes to the "Custom Resource" and create it.

```
kubectl create -f k8sdc_cr.yaml
```

A template of the `k8sdc_cr.yaml` file is shown here:

```
apiVersion: ibmcloudappmgmt.com/v1alpha1
kind: K8sDC
metadata:
  name: icamklust-k8sdc-dc
  namespace: multicluster-endpoint
spec:
  arch: ""
  clusterName: my_cluster_name
  collectEvents: "on"
  collectMetrics: "off"
  collectResources: "off"
  global:
    environmentSize: size0
    imagePullPolicy: IfNotPresent
  ibmAgentConfigSecret: dc-secret
  ibmAgentHTTPSSecret: ""
  image:
    repository: my_cluster_ca_domain:my_docker_registry_port/my_namespace
    imageNamePrefix: ""
  imageTag: APM_201909192049
  rbac:
    create: true
  replicas: 1
```

### c) Create the service account for the ICAM klusterlet:

```
kubectl create -f service_account.yaml
```

A template of the `service_account.yaml` file is shown here:

```
apiVersion: v1
kind: ServiceAccount
metadata:
  name: icamklust
  namespace: my_namespace
```



- d) 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 my_namespace
```

- e) 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=my_namespace:icamklust -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).

- f) Create an empty dc-secret:

```
kubectl create -f dc-secret.yaml
```

A template of the `dc-secret.yaml` file is shown here:

```
apiVersion: v1
kind: Secret
metadata:
  name: dc-secret
type: Opaque
data:
  APM_TENANT_ID: ''
  IBM_APM_SERVER_INGRESS_URL: ''
  IBM_CEM_APIKEY_PASSWORD: ''
  IBM_CEM_EVENT_INGRESS_URL: ''
```

- g) Create the k8sdc-operator deployment:

```
kubectl create -f k8sdc-operator.yaml
```

A template of the `k8sdc-operator.yaml` file is shown here:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: k8sdc-operator
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: REPLACE_IMAGE
          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"
```

- h) Create the reloader deployment:

```
kubectl create -f icam-reloader.yaml
```

A template of the `icam-reloader.yaml` file is shown here:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: icam-reloader
spec:
  replicas: 1
  revisionHistoryLimit: 2
  selector:
    matchLabels:
      name: icam-reloader
  template:
    metadata:
      labels:
        name: icam-reloader
    spec:
      containers:
        - env:
            - name: KUBERNETES_NAMESPACE
              valueFrom:
                fieldRef:
                  apiVersion: v1
                  fieldPath: metadata.namespace
            # Replace this with the built image name
            image: REPLACE_IMAGE
            imagePullPolicy: Always
            name: icam-reloader
```

i) Create the agentoperator deployment:

```
kubectl create -f agentoperator.yaml
```

A template of the `agentoperator.yaml` file is shown here:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: agentoperator
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: REPLACE_IMAGE
          imagePullPolicy: Always
          command:
            - agentoperator
          args:
            - "--namespace=my-remote-namespace"
            - "--cluster-name=my-cluster-name"
            - "--cluster-labels=my-cluster-labels"
          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: my-mcm-klusterlet-hub-kubeconfig
```

Where:

- *my-remote-namespace* is the remote namespace in hub cluster you want to watch
- *my-cluster-labels* are the set of labels that are used in Prometheus. For example: "cloud=IBM,vendor=ICP,environment=Dev,region=US,datacenter=toronto,owner=marketing"
- *my-mcm-klusterlet-hub-kubeconfig* is the secret name that is created by the IBM Multicloud Manager klusterlet, which includes kubeconfig to connect to hub cluster. For example: "multicluster-endpoint-hub-kubeconfig".

Validate the deployments that you created in the previous steps:

7. After the installation script is finished, run the following commands:

```
kubectl get deployment k8sdc-operator --namespace=my_namespace
kubectl get deployment icam-reloader --namespace=my_namespace
kubectl get deployment agentoperator --namespace=my_namespace
```

### What to do next

1. Update the IBM Multicloud Manager team for Cloud App Management. The namespace corresponding to any new managed cluster must be added as a resource to the Cloud App Management team. On the IBM Cloud Private console, click **Manage > Identity and Access > Teams**. For more information, see [Adding the Helm repository and namespace to a team](#). Click the **Event Management** menu item to update the ICAM klusterlet with the correct webhook for Cloud App Management events.
2. 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 Multicloud Manager”](#) on page 89.

## Deploying agents and data collectors for IBM Multicloud Manager

After you deploy the ICAM klusterlet (icam-clouddc-klusterlet) for IBM Multicloud Manager, you can install and configure agents and data collectors.

### Before you begin

The ICAM klusterlet for IBM Multicloud Manager must be deployed. For more information, see [“Deploying the ICAM klusterlet for IBM Multicloud Manager ”](#) on page 83.

### About this task

Kubernetes data collector is auto-configured when you deploy the ICAM klusterlet for IBM Multicloud Manager. The other agents and data collectors still need manual configuration. The steps to deploy these agents and data collectors for IBM Multicloud Manager are similar to the instructions in IBM Cloud Private environment.

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 **IBM Multicloud Manager** dashboard.
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 9, “Deploying ICAM Agents,” on page 131.](#)
- [Chapter 11, “Deploying ICAM Data Collectors,” on page 427.](#)
- 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 12, “Unified Agent,” on page 499.](#)

## Uninstalling IBM Cloud App Management for IBM Multicloud Manager

If you no longer want Cloud App Management installed on your system with IBM Multicloud Manager, you can delete the helm deployment.

### Before you begin

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 125.](#)

### Procedure

Complete the following steps to uninstall the Cloud App Management server:

1. Log in to your IBM Cloud Private cluster using the IBM Cloud Private CLI:

```
cloudctl login -a https://my_cluster_CA_domain:8443
--skip-ssl-validation
```

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 *my\_cluster\_name.icp* where *my\_cluster\_name* is the name you defined for your cluster. The default *my\_cluster\_name* is *mycluster*.

2. Find the Helm chart that you want to uninstall from the list:

```
helm list --tls | grep ibm-cloud-appmgmt
```

3. 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.

4. 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
```

where *my\_namespace* is the namespace that the IBM Passport Advantage Archive (PPA) file is loaded to.

5. Delete secrets and the cluster image policy:

```
kubectl delete secrets --selector release=my_release_name --namespace my_namespace
kubectl delete clusterimagepolicy --selector release=my_release_name --namespace
my_namespace
```

6. Optional: Back up the data on the persistent storage directories that you created on the worker nodes.
7. Delete persistent volumes:

```
kubectl delete pv --selector release=my_release_name
```

8. Optional: You can safely remove the data from the persistent storage directories that you created on the worker nodes.
9. 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

The Cloud App Management server helm chart is uninstalled. The storage configuration that was required for the installation was also deleted.

## Installing IBM Cloud App Management on Red Hat OpenShift

These steps explain how to install the IBM Cloud App Management server on Red Hat OpenShift.

### Before you begin

- Cloud Pak Foundation V3.2.1 is required. Download the Cloud Pak Foundation Docker package: `ibm-cloud-private-rhos-3.2.1.tar.gz` from the [IBM Passport Advantage](#) website.
- Red Hat OpenShift V3.11 is required. You must also install the OpenShift Container Platform CLI. For more information, see [Get Started with the CLI](#).
- Install the Helm CLI commands. For instructions, see [Installing the Helm CLI \(helm\)](#).
- Install the Kubernetes command line tool, `kubectl`, and configure access to your cluster. See [Accessing your cluster from the Kubernetes CLI \(kubectl\)](#).

### Procedure

1. Locate the Cloud App Management server installation image file `app_mgmt_server_2019.3.0.tar.gz` (part number: CC3FIEN) on [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 the [“Part numbers” on page 53](#).

**Note:** The Cloud App Management server with IBM Cloud Pak for Multicloud Management must be installed in the `kube-system` namespace.

2. As the IBM Cloud Private administrator, log in to the management console and select your namespace. Run the following command:

```
cloudctl login -a my_cluster_URL -n kube-system -skip-ssl-validation
```

Where `my_cluster_URL` is the IBM Cloud Private name that you defined for your cluster such as `https://cluster_address:443`. For any ‘masterIP’ reference, use the `cluster_address` value.

**Note:** Run the following command to get your cluster information before you start.

```
kubectl get configmap -n kube-public -o yaml
```

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)
```

5. Create an install directory and download the IBM Passport Advantage file: `app_mgmt_server_2019.3.0.tar.gz` to the install directory. Change to the install directory:

```
mkdir -p install
cd install
```

6. Extract the Helm charts from the Passport Advantage Archive (PPA) file into the `install_dir` directory. Extract the (PPA) file by running the following example commands:

```
cd install_dir
tar -xvf ./ppa_file charts
tar -xvf ./charts/ibm-cloud-appmgmt-prod-1.5.0.tgz
```

Where:

- `ppa_file` is the compressed Cloud App Management PPA installation image file: `app_mgmt_server_2019.3.0.tar.gz` file. The `charts` value is required to ensure the `tar` command extracts only the `charts` directory from the `ppa_file`. Otherwise, all the images are extracted, which might cause space issues.

7. Load the PPA file into IBM Cloud Private Docker registry:

```
cloudctl catalog load-archive --archive ./ppa_file --registry $(oc registry info)/kube-system
```

8. Change directory to the `ibm-cloud-appmgmt-prod` directory.

```
cd install_dir/ibm-cloud-appmgmt-prod
```

9. Create the Cloud App Management ingress TLS and client secrets:

```
./ibm_cloud_pak/pak_extensions/lib/make-ca-cert-icam.sh my_ProxyHostName my_release_name kube-system
```

Where `my_release_name` is the name of the Cloud App Management Helm Chart, for example: `ibmcloudappmgmt`.

10. Apply the OpenShift security context constraints into the namespace where you are installing Cloud App Management.

**Note:** This chart requires a `SecurityContextConstraints` to be bound to the target namespace before you install Cloud App Management. To meet this requirement, cluster and namespace preinstallation and postinstallation steps might need to be completed. The predefined `SecurityContextConstraints` name: `ibm-restricted-scc` is verified for this chart. If your target namespace is bound to this `SecurityContextConstraints` resource, you can start to install the chart. This chart also defines a custom `Constraints`, which can be used to control the permissions and capabilities that are needed to deploy this chart. You can enable this custom `SecurityContextConstraints` resource by using the following scripts.

- a) Run the following setup script as a cluster administrator. The postinstallation instructions are included in this script.

```
cd ibm_cloud_pak/pak_extensions/pre-install/clusterAdministration
./createSecurityClusterPrereqs.sh
```

- b) Run the following setup script as a team administrator. The namespace instructions are included in this script.

```
cd ibm_cloud_pak/pak_extensions/pre-install/namespaceAdministration
./createSecurityNamespacePrereqs.sh kube-system
```

- c) Change directories back to the `install_dir/ibm-cloud-appmgmt-prod` directory. For example: `cd /install/ibm-cloud-appmgmt-prod`

11. Update the following global values in the Cloud App Management `values.yaml` file.

```

global:
  masterIP: 'icp-console.apps.masternode.ibm.com'
  masterPort: 443
  masterCA: 'icam-cluster-ca-cert'
  ingress:
    domain: 'icp-proxy.apps.masternode.ibm.com'
    port: 443

```

- Remove port 80 from the nginx-ingress service:

```
kubectl edit services -n kube-system nginx-ingress
```

Remove the lines that have a line through them below.

```

apiVersion: v1
kind: Service
metadata:
  annotations:
    kubectll.kubernetes.io/last-applied-configuration: |
      {"apiVersion":"v1","kind":"Service","metadata":{"annotations":{},"name":"nginx-
      ingress","namespace":"kube-system"},"spec":{"ports":
      [{"name":"http","port":80,"protocol":"TCP","targetPort":80},
      {"name":"https","port":443,"protocol":"TCP","targetPort":443}], "selector":{"app":"nginx-
      ingress-controller"},"type":"ClusterIP"}}
    creationTimestamp: 2019-05-17T18:44:00Z
    name: nginx-ingress
    namespace: kube-system
    resourceVersion: "54038"
    selfLink: /api/v1/namespaces/kube-system/services/nginx-ingress
    uid: bc20e4de-78d3-11e9-9373-00163e01c818
spec:
  clusterIP: 172.30.79.16
  ports:
  - name: http
  - port: 80
  - protocol: TCP
  - targetPort: 80
  - name: https
    port: 443
    protocol: TCP
    targetPort: 443
  selector:
    app: nginx-ingress-controller
  sessionAffinity: None
  type: ClusterIP
status:
  loadBalancer: {}

```

- Create local directories on each of the IBM Cloud Private worker nodes. Record the IP address of each worker node where the directories are created. You need these values when you are selecting your storage class and preparing persistent volumes in the next step in this procedure. Review the optimization performance topics. For more information, see [“Optimizing performance” on page 107](#).

**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.

- /k8s/data/cassandra - Cassandra persistent storage

```
mkdir -p /k8s/data/cassandra
```

- /k8s/data/zookeeper - ZooKeeper persistent storage

```
mkdir -p /k8s/data/zookeeper
```

- /k8s/data/kafka - Kafka persistent storage

```
mkdir -p /k8s/data/kafka
```

- /k8s/data/couchdb - CouchDB persistent storage

```
mkdir -p /k8s/data/couchdb
```

- /k8s/data/datalayer - Datalayer persistent storage

```
mkdir -p /k8s/data/datalayer
```

- /k8s/data/elasticsearch - Elasticsearch persistent storage

```
mkdir -p /k8s/data/elasticsearch
```

14. Select your storage class and prepare the persistent volumes for each of the following Cloud App Management services by running the `prepare-pv.sh` script. Local storage is recommended. The services are:

- Cassandra
- Kafka
- ZooKeeper
- CouchDB
- Data layer
- Elasticsearch

**Note:** For additional information, see the "Storage" section in the **Overview** tab by selecting the IBM-cloud-appmgmt-prod tile from the catalog. For information about how to configure persistent storage, see the following topics: [“Planning hardware and sizing”](#) on page 58, [Understanding Kubernetes storage](#), [Planning a storage solution](#), [Planning persistent storage](#), and [Storage options in IBM Cloud Private](#)

**Note:** Before you run the `prepare-pv.sh` script, determine the addresses of the worker nodes where the persistent storage directories are located for each service by running the following command:

```
kubectl get nodes
```

The addresses can be the IP or hostname depending on what IBM Cloud Private is using.

- For local storage, specify the **--local** parameter. Local storage is recommended.
- For vSphere, specify the **--vSphere** parameter.
- Choose from predefined `--size0_amd64`, `--size0_ppc64le`, `--size1_amd64`, or `--size1_ppc64le` storage values. For more information about sizes, see [“Planning hardware and sizing”](#) on page 58.

```
ibm_cloud_pak/pak_extensions/prepare-pv.sh
--size0_amd64      #Install as size0 on amd64 (minimum resource requirements) - if
omitted, specify each size in the size parameters below
--size0_ppc64le   #Install as size0 on ppc64le (minimum resource requirements) - if
omitted, specify each size in the size parameters below
--size1_amd64     #Install as size1 on amd64 (standard resource requirements) - For
HA use --size1_amd64 - if omitted, specify each size in the size parameters below
--size1_ppc64le   #Install as size1 on ppc64le (standard resource requirements) -
For HA use --size1_ppc64le - if omitted, specify each size in the size parameters below
--releasename     #installation name, such as ibmcloudappmgmt, as defined in your
cluster image policy.

#Required flags for local storage
--local          #Use local persistent volume storage
#For high availability, list the nodes in quotes, separated by spaces
#If IBM Cloud
Private uses the IP address instead of the hostname, the IP address is needed here.
--CassandraNodes #IP or hostname of the node/s with the persistent storage local
directory for the Cassandra service.
--ZookeeperNodes #IP or hostname of the node/s with the persistent storage local
directory for the Zookeeper service.
--KafkaNodes     #IP or hostname of the node/s with the persistent storage local
directory for the Kafka service.
--CouchDBNodes   #IP or hostname of the node/s with the persistent storage local
directory for the CouchDB service.
--DatalayerNodes #IP or hostname of the node/s with the persistent storage local
directory for the Datalayer service.
--ElasticsearchNodes #IP or hostname of the node/s with the persistent storage local
```



directory for the Elasticsearch service.

```
#Optional storage directory paths for local storage
--CassandraDir #the local system directory for Cassandra (default is /k8s/data/cassandra)
--CassandraBackupDir #the local system backup directory for Cassandra (default is /k8s/data/cassandra_backup)
--KafkaDir #the local system directory for Kafka (default is /k8s/data/kafka)
--ZookeeperDir #the local system directory for Zookeeper (default is /k8s/data/zookeeper)
--CouchDBDir #the local system directory for CouchDB (default is /k8s/data/couchdb)
--DatalayerDir #the local system directory for Datalayer (default is /k8s/data/datalayer)
--ElasticsearchDir #the local system directory for Elasticsearch (default is /k8s/data/elasticsearch)

#Optional storage class name flags for local storage:
--CassandraClass #the storage class name for Cassandra (default is <release_name>-local-storage-cassandra)
--CassandraBackupClass ##the storage class name for Cassandra backups (default is <release_name>-local-storage-cassandrabackup)
--KafkaClass #the storage class name for Kafka (default is <release_name>-local-storage-kafka)
--ZookeeperClass #the storage class name for Zookeeper (default is <release_name>-local-storage-zookeeper)
--CouchDBClass #the storage class name for CouchDB (default is <release_name>-local-storage-couchdb)
--DatalayerClass #the storage class name for Datalayer (default is <release_name>-local-storage-datapayer)
--ElasticsearchClass #the 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:
#Persistent volumes are measured in bytes.
#When you specify the size, you must use only the number and prefix, for example, 50 Gi. If you specify bytes, such as 50 GiB, the installation fails.
--CassandraSize #the size of persistent volume for Cassandra (default size0 is 50 Gi)
--CassandraBackupSize #the size of persistent volume for Cassandra backups (default size0 is 50 Gi)
--KafkaSize #the size of persistent volume for Kafka (default size0 is 10 Gi)
--ZookeeperSize #the size of persistent volume for Zookeeper (default size0 is 1 Gi)
--CouchDBSize #the size of persistent volume for CouchDB (default size0 is 1 Gi)
--DatalayerSize #the size of persistent volume for Datalayer (default size0 is 1 Gi)
--ElasticsearchSize #the size of persistent volume for Elasticsearch (default size0 is 1 Gi)
```

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"
```

All directories for a statefulset type defined by `prepare-pv.sh` are in the same directory, for example: `/k8s/data/cassandra/`. If you want to use different directories, after you run `prepare-pv.sh` you can customize the `.yaml` file. Complete any customization to the `yaml` file before you run the `kubectl create` command. By default the `yaml` files are generated in: `ibm-cloud-appmgmt-prod/ibm_cloud_pak/yaml/`

15. Configure the Elasticsearch `vm.max_map_count` parameter:

- a) For Elasticsearch, you must set a kernel parameter to run normally for all worker nodes where it runs. These nodes are identified when you configure the persistent storage. You need to set `vm.max_map_count` to a minimum value of 1048575. Set the parameter with `sysctl` to ensure that the change takes effect immediately. For example:

```
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
```

16. Run the `pre-install.sh` script with the following parameters, note the following important choices first:

- Use the `--https` parameter for IBM Cloud App Management, Base or IBM Cloud App Management, Advanced. Switching between HTTP and HTTPS communication is not supported after installation.
- Use the `--advanced` parameter to install IBM Cloud App Management, Advanced. HTTPS communication is always enabled.
- Use the `--masterPort` for the IBM Cloud Private master. Default port for OpenShift is 443.
- Use the `--repositoryPort` for the image repository. The default port for OpenShift is 5000.
- To ensure that your environment is secure, provide a valid `--cassandraUsername` instead of the default "cassandra". When the `--cassandraUsername` parameter is passed to `pre-install.sh`, the script prompts you to enter a unique secure password to use for the Cassandra cluster rather than accepting the default value.

Usage `./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                Accept 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 <IP|FQDN> IP address or fully qualified domain name (FQDN)
for the ICP Master. You must use the address that you used to login to the ICP console.
In a highly available environment, this would be
the FQDN of the HAProxy or load balancer for ICP.
--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
--cassandraUsername <string> The username Cassandra will use. If left unset,
the default cassandra credentials will be used.

*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)
--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)

*Optional - Other
--masterPort <int> The port of the ICP master. The default port for
OpenShift is 443.
--repositoryPort <int> The port of the image repository. The default
port for OpenShift is 5000.
--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)

```

**Note:** The `--cassandraUsername` parameter should be set. Specify a value or the default user name "cassandra" is used.

Example scenarios:

Scenario 1: No High Availability

```

ibm_cloud_pak/pak_extensions/pre-install.sh --accept --releasename ibmcloudappgmt --
namespace default --masterAddress icp-master.mydomain.com --proxyIP 0.1.2.3
--proxyFQDN icp-proxy.mydomain.com --clustercadomain mycluster.icp --advanced --
cassandraclustersize 1 --kafkaclustersize 1 --zookeeperclustersize 1 --couchdbclustersize 1
--datalayerclustersize 1 --elasticsearchclustersize 1 --minreplicashpas 1 --
maxreplicashpas 3 --cassandraUsername myCassandraUser

```

Scenario 2: High Availability with 3 Cassandra and 2 of each deployment (HPA minimums)

```

ibm_cloud_pak/pak_extensions/pre-install.sh --accept --releasename ibmcloudappgmt --
namespace default --masterAddress icp-master.mydomain.com --proxyIP 0.1.2.3
--proxyFQDN icp-proxy.mydomain.com --clustercadomain mycluster.icp --advanced --
cassandraclustersize 3 --kafkaclustersize 3 --zookeeperclustersize 3 --couchdbclustersize 3
--datalayerclustersize 3 --elasticsearchclustersize 3 --minreplicashpas 2 --
maxreplicashpas 3 --cassandraUsername myCassandraUser

```

Scenario 3: High Availability with 6 Cassandra and 2 of each deployment (HPA minimums)

```

ibm_cloud_pak/pak_extensions/pre-install.sh --accept --releasename ibmcloudappgmt --
namespace default --masterAddress icp-master.mydomain.com --proxyIP 0.1.2.3
--proxyFQDN icp-proxy.mydomain.com --clustercadomain mycluster.icp --advanced --
cassandraclustersize 6 --kafkaclustersize 3 --zookeeperclustersize 3 --couchdbclustersize 3
--datalayerclustersize 3 --elasticsearchclustersize 3 --minreplicashpas 2 --
maxreplicashpas 3 --cassandraUsername myCassandraUser

```

- Optional: Review and update the Helm chart configuration `my_release_name.values.yaml`. For more information, see [“Configuring the Helm charts”](#) on page 119.

18. Optional: If you want to either change the raw metric retention period from the default 8 days or enable metric summarization, add one (or both) of the following `--set` forms to the **helm install** command in step “19” on page 98:

```
--set global.metric.retention.rawMaxDays=2
```

where 2 represents the number of days to retain and can be a whole number from 2 to 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 607.

This 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-1.5.0.tgz --tls`.

19. Deploy the Cloud App Management server Helm chart using the Helm CLI, by run the following command:

```
helm install --name my_release_name --values my_release_name.values.yaml my_install_dir/ibm-cloud-appgmt-prod-1.5.0.tgz --tls
```

Where `my_install_dir` is the directory where you extracted the Helm chart.

20. Run the `post-install-setup.sh` script to complete administrative tasks necessary to access the IBM Cloud App Management dashboard. Use the following parameters, note the following important choices first:

- **--releaseName --instanceName --namespace**
- Use: **--advanced** for Advanced offering.

```
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"
--releaseName ${default_release} --instanceName ${instance_name} --namespace ${namespace}"
"example: for Advanced offering"
--releaseName ${default_release} --instanceName ${instance_name} --namespace ${namespace} --
advanced"
```

Example:

```
my_install_dir/ibm-cloud-appgmt-prod/ibm_cloud_pak/pak_extensions>/post-install-setup.sh --
releaseName ibmcloudappgmt

--namespace icam --instanceName ibmcloudappgmt --advanced
```

## Results

The Cloud App Management server is successfully installed.

## What to do next

1. Create a service instance. For more information, see “Creating your service instance” on page 123.
2. Start the service instance and access the Cloud App Management console. For more information, see “Starting the Cloud App Management UI” on page 124.

## Installing IBM Cloud App Management 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 6, “Planning your deployment,” on page 51](#).

1. Deploy IBM Cloud Private. For more information, see [Chapter 7, “Deploying IBM Cloud Private,” 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 107](#)
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 100](#).
3. **Optional:** Configure a custom certificate. For more information, see [“Configuring a custom server certificate” on page 104](#).
4. Install the Cloud App Management server. This involves the following sub steps, which are described in detail in the [“Installing on IBM Cloud Private” on page 112](#) topic:

Prepare your IBM Cloud Private nodes by creating local directories on each IBM Cloud Private worker node.

Download the Cloud App Management installation image file.

Extract the Helm charts from the installation file and load the installation file into the Docker repository.

Prepare persistent volumes by running the `prepare-pv.sh` script. Local storage is recommended as it allows you to lock storage to the node and the local directory.

Define your installation parameters and optionally you may further configure the Helm chart. For more information, see [“Configuring the Helm charts” on page 119](#).

Install the Cloud App Management server.

For more information, see [“Installing on IBM Cloud Private” on page 112](#).

5. **Optional:** Move to a custom namespace. For more information, see [“Moving to a custom namespace” on page 121](#).
6. **Optional:** Validate the Cloud App Management server deployment is successful. For more information, see [“Validating the Cloud App Management server deployment” on page 122](#).
7. Complete the post-installation task of creating your service instance. For more information, see [“Creating your service instance” on page 123](#).
8. Access the Cloud App Management console. For more information, see [“Starting the Cloud App Management UI” on page 124](#).

## Downloading the Cloud App Management Passport Advantage Archive (PPA) file

Before you can deploy the Cloud App Management server, you must download and extract the Cloud App Management installation image file from [IBM Passport Advantage](#).

### Procedure

Go to the [IBM Passport Advantage](#) website. Sign in and identify the Passport Advantage Archive (PPA) file that you want to download for the Cloud App Management product, such as `app_mgmt_server_2019.3.0.tar.gz`. For more information, see [“Part numbers”](#) on page 53.

### Results

The PPA file is downloaded.

### What to do next

Extract the Helm chart from the PPA file and install the Cloud App Management server. For instructions, see [“Installing on IBM Cloud Private”](#) on page 112.

## 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


HPA provides scalability for stateless services as described here: [“Scaling stateless and stateful services”](#) on page 63. The HPAs have a minimum and maximum number of replicas they can scale to. This value is determined by the `minReplicasHPAs` and `maxReplicasHPAs` global parameters, which are specified when you run `pre-install.sh`. 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 `pre-install.sh` 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 63.

### 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 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 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 keyspace 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. For more information see, step “10” on page 116 in the *Installing on IBM Cloud Private* topic.

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 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.

The default replication for metric data in Kafka (when deployed with 2 or more Kafka) has been changed from 3 to 2 to reduce the backend overhead of replicating this data, as it is the largest set of 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 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 on IBM Cloud Private” on page 112](#) 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 the [Table 9 on page 59](#) table in the *Planning hardware and sizing* topic.



2. Proceed with your installation as described in [“Installing on IBM Cloud Private”](#) on page 112. At the following points, during your installation, you must specify parameters based on your resiliency planning:

a. In step [“8”](#) on page 114, 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. In step [“10”](#) on page 116, 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 63.

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
--clustercadomain 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 123 topic and also in the [“Freeing up resources for large Pod scheduling”](#) on page 650 topic.

For information, about upgrading to a highly available environment, see [“Upgrade to a high-availability environment”](#) on page 630.

## 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 104](#).

**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 in 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 132](#)
- [“Connecting IBM Tivoli Monitoring agents to Cloud App Management server” on page 518](#)
- [“Connecting Cloud APM agents to Cloud App Management server” on page 532](#)
- [Chapter 11, “Deploying ICAM Data Collectors,” on page 427](#)

## 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 `ibmcloudappmgmt`. 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 ibmcloudappmgmt-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 `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 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 132](#)
- [“Connecting IBM Tivoli Monitoring agents to Cloud App Management server” on page 518](#)
- [“Connecting Cloud APM agents to Cloud App Management server” on page 532](#)
- [Chapter 11, “Deploying ICAM Data Collectors,” on page 427](#)

#### 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 `ibmcloudapmgmt`. 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 ibmcloudappmgmt-ingress-client --namespace=my_namespace -o yaml  
> ibmcloudappmgmt-ingress-client.backup.yaml
```

c) Replace the current secret with your new certificate by running the following command:

```
kubectl create secret generic ibmcloudappmgmt-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/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 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.services` disk settings to make the readahead settings persistent. These steps need to be performed on each VM running Cassandra.

### Manually Setting Readahead on Existing Drive or Volume

In order to set the readahead, use the `blockdev` command with the internal kernel device name (KNAME). The KNAME of the device to modify can be found by running the command:

Note, in this example we are running Cassandra on `dm-2`. Your KNAME may be different based on your system settings.

```
lsblk --output NAME,KNAME,TYPE,MAJ:MIN,FSTYPE,SIZE,RA,MOUNTPOINT,LABEL
```

```
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    4096
└─vg_sdc-lv_cassandra dm-2 lvm 253:2    xfs    2T    4096 /k8s/data/cassandra cassandra
```

The `--setra` will be in number of blocks, so a readahead of 16 with size of 512 bytes results in an 8KB readahead.

```
blockdev --setra 16 <device>
```

For example:

```
blockdev --setra 16 /dev/dm-2
```

Verify the readahead settings:

```
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     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, the Cassandra docker container will need to be restarted in order to use the new readahead values. This can be done as an IBM Cloud Private admin, either through the IBM Cloud Private UI or `kubectl` command.

## Procedure

**Modify tuned.service Disk Settings To Make Readahead Persistent** 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. 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*.

1. Format a blank drive for the Cassandra data to be stored as described in [“Configuring the disk drives for services”](#) on page 110.
2. 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.

3. 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/
```

4. 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
```

5. Reload the tuned profile:

```
tuned-adm profile virtual-guest
```

6. Verify the new readahead setting:

```
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  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.

### 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 5 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 4 services (Zookeeper, Kafka, CouchDB, and Datalayer) 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 4 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
```

#### 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
```

#### 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
```

#### 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
```

#### 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
```

#### 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	data1

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 107](#)

## Installing 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

When you install the Cloud App Management server, use IBM Cloud Private V3.2.1. The Cloud App Management server V2019.3.0 runs on IBM Cloud Private V3.2.1.

**Note:** 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
```

Review the optimization for performance topics. For more information, see [“Optimizing performance” on page 107](#)

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 104](#).

Familiarize yourself with the IBM Cloud App Management offerings. For more information, see [“Offerings” on page 14](#).

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.

## Procedure

Complete the following steps as an IBM Cloud Private cluster administrator:

1. Optional: As the IBM Cloud Private administrator, log in to the management console and select your namespace. Run the following command:

```
cloudctl login -a my_cluster_URL --skip-ssl-validation
```

Where *my\_cluster\_URL* is the IBM Cloud Private name you defined for your cluster such as `https://<cluster_address>:443`. For any 'masterIP' reference, use the `cluster_address` value.

**Note:** Run the following command to get your `cluster_address` and port:

```
kubectl get configmap -n kube-public -o yaml
```

2. Create an installation *install\_dir* directory. You can specify any name for the directory. In the example, `install` is used. Go to the IBM Passport Advantage website. Sign in and download the following Passport Advantage Archive (PPA) file `app_mgmt_server_2019.3.0.tar.gz` into the *install\_dir* directory. For more information, see [“Part numbers” on page 53](#). Change directory to the *install\_dir* directory. For example,

```
mkdir -p install
```

```
cd install
```

**Note:** The Cloud App Management server installation image (PPA file) is large. Be sure that you have enough free space to store the file.

3. Extract the Helm charts from the Cloud App Management server installation image file by running the following commands from the `install` directory where you saved the `app_mgmt_server_2019.3.0.tar.gz` file:

```
tar -xvf app_mgmt_server_2019.3.0.tar.gz charts
```

```
tar -xvf charts/ibm-cloud-appmgmt-prod-1.5.0.tgz
```

**Note:** The `charts` value is required and ensures the `tar` command extracts only the `charts` directory from the PPA file. Otherwise, all the images are extracted leading to potential space issues.

4. Load the PPA file into Docker.
  - a) Log in to the Docker private image 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`.

- b) Load the Cloud App Management PPA file into the IBM Cloud Private local repository by running the following command:

```
cloudctl catalog load-archive --archive app_mgmt_server_2019.3.0.tar.gz  
[--registry my_cluster_CA_domain:8500] [--repo my_helm_repo_name]
```

Where *my\_helm\_repo\_name* is the name of the target Helm repository. Run the **cloudctl catalog repos** command to get a list of repositories.

**Note:** Unless you specify an `imagePullSecret`, you can access this image from only the namespace that hosts it. If your environment does not use `mycluster.icp:8500` as the registry parameter, you must log in to your registry and specify this parameter. Similarly, you can specify the `--repo` parameter in your environment.

5. Change directory to the `ibm-cloud-appmgmt-prod` directory.

```
cd ibm-cloud-appmgmt-prod
```

## 6. Elasticsearch vm.max\_map\_count requirement.

Elasticsearch requires you to set a kernel parameter to run normally. This step needs to be completed on all worker nodes where Cloud App Management is installed. These nodes are determined when you configure the persistent storage. You need to set `vm.max_map_count` to a value of at least 1048575. Set the parameter with `sysctl` to ensure that the change takes effect immediately:

```
sysctl -w vm.max_map_count=1048575
```

You can also set the parameter in `/etc/sysctl.conf` to preserve the changes after a node restart by adding:

```
vm.max_map_count=1048575
```

## 7. Create local directories as follows on each IBM Cloud Private worker node where the persistent storage is used, and record the IP address of each worker node where the directories are created. You will need these values in the next step. Review the optimization performance topics. For more information, see [“Optimizing performance”](#) on page 107.

**Note:** For best performance, create the Cassandra persistent volume on it's own disk and dedicated worker node. Avoid creating the other persistent volumes on the same worker node and the same PV with Cassandra.

- `/k8s/data/cassandra` - Cassandra persistent storage

```
mkdir -p /k8s/data/cassandra
```

- `/k8s/data/zookeeper` - Zookeeper persistent storage

```
mkdir -p /k8s/data/zookeeper
```

- `/k8s/data/kafka` - Kafka persistent storage

```
mkdir -p /k8s/data/kafka
```

- `/k8s/data/couchdb` - CouchDB persistent storage

```
mkdir -p /k8s/data/couchdb
```

- `/k8s/data/datalayer` - Datalayer persistent storage

```
mkdir -p /k8s/data/datalayer
```

- `/k8s/data/elasticsearch` - Elasticsearch persistent storage

```
mkdir -p /k8s/data/elasticsearch
```

## 8. Select your storage class and prepare the persistent volumes by running the `prepare-pv.sh` script. Local storage is recommended as it allows you to lock storage to the node and the local directory.

**Note:** Before you run the `prepare-pv.sh` script. Determine the addresses of the worker nodes where the persistent storage directories are located for each service. Run the following command:

```
kubectl get nodes
```

The addresses can be the IP or hostname based on what IBM Cloud Private is using.

**Note:** More information about storage classes is available in the [Chapter 6, “Planning your deployment,”](#) on page 51 topic.

- For local storage, specify the `--local` parameter. Local storage is best practice.
- For vSphere, specify the `--vsphere` parameter.

- Choose from predefined `--size0_amd64`, `--size0_ppc64le`, `--size1_amd64`, or `--size1_ppc64le` storage values. For more information about sizes, see [“Planning hardware and sizing”](#) on page 58.

```

ibm_cloud_pak/pak_extensions/prepare-pv.sh
--size0_amd64      #Install as size0 on amd64 (minimum resource requirements) - if
omitted, specify each size in the size parameters below
--size0_ppc64le   #Install as size0 on ppc64le (minimum resource requirements) - if
omitted, specify each size in the size parameters below
--size1_amd64     #Install as size1 on amd64 (standard resource requirements) - For
HA use --size1_amd64 - if omitted, specify each size in the size parameters below
--size1_ppc64le  #Install as size1 on ppc64le (standard resource requirements) -
For HA use --size1_ppc64le - if omitted, specify each size in the size parameters below
--releaseName     #Installation name, such as ibmcloudappmgmt, as defined in your
cluster image policy.

    #Required flags for local storage
--local           #Use local persistent volume storage
    #For high availability, list the nodes in quotes, separated by spaces
    #If IBM Cloud
Private uses the IP address instead of the hostname, the IP address is needed here.
--CassandraNodes #IP or hostname of the node/s with the persistent storage local
directory for the Cassandra service.
--ZookeeperNodes #IP or hostname of the node/s with the persistent storage local
directory for the Zookeeper service.
--KafkaNodes     #IP or hostname of the node/s with the persistent storage local
directory for the Kafka service.
--CouchDBNodes   #IP or hostname of the node/s with the persistent storage local
directory for the CouchDB service.
--DatalayerNodes #IP or hostname of the node/s with the persistent storage local
directory for the Datalayer service.
--ElasticsearchNodes #IP or hostname of the node/s with the persistent storage local
directory for the Elasticsearch service.

    #Optional storage directory paths for local storage
--CassandraDir   #the local system directory for Cassandra (default is /k8s/data/
cassandra)
--CassandraBackupDir #the local system backup directory for Cassandra (default is /k8s/
data/cassandra_backup)
--KafkaDir       #the local system directory for Kafka (default is /k8s/data/kafka)
--ZookeeperDir   #the local system directory for Zookeeper (default is /k8s/data/
zookeeper)
--CouchDBDir     #the local system directory for CouchDB (default is /k8s/data/
couchdb)
--DatalayerDir   #the local system directory for Datalayer (default is /k8s/data/
datalayer)
--ElasticsearchDir #the local system directory for Elasticsearch (default is /k8s/
data/elasticsearch)

    #Optional storage class name flags for local storage:
--CassandraClass #the storage class name for Cassandra (default is <release_name>-
local-storage-cassandra)
--CassandraBackupClass ##the storage class name for Cassandra backups (default is
<release_name>-local-storage-cassandrabackup)
--KafkaClass     #the storage class name for Kafka (default is <release_name>-
local-storage-kafka)
--ZookeeperClass #the storage class name for Zookeeper (default is <release_name>-
local-storage-zookeeper)
--CouchDBClass   #the storage class name for CouchDB (default is <release_name>-
local-storage-couchdb)
--DatalayerClass #the storage class name for Datalayer (default is <release_name>-
local-storage-datapayer)
--ElasticsearchClass #the 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:
    #Persistent volumes are measured in bytes.
    #When you specify the size, you must use only the number and prefix, for example, 50
Gi. If you specify bytes, such as 50 GiB, the installation fails.
--CassandraSize  #the size of persistent volume for Cassandra (default size0 is 50
Gi)
--CassandraBackupSize #the size of persistent volume for Cassandra backups (default
size0 is 50 Gi)
--KafkaSize      #the size of persistent volume for Kafka (default size0 is 10 Gi)
--ZookeeperSize  #the size of persistent volume for Zookeeper (default size0 is 1
Gi)

```

```

--CouchDBSize          #the size of persistent volume for CouchDB (default size0 is 1 Gi)
--DataLayerSize        #the size of persistent volume for DataLayer (default size0 is 1
Gi)
--ElasticsearchSize    #the size of persistent volume for Elasticsearch (default size0
is 1 Gi)

```

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"

```

All of the directories for a statefulset type defined by `prepare-pv.sh` are in the same directory, for example, `/k8s/data/cassandra/`. If you want to use different directories, after running `prepare-pv.sh` you can customize the YAML file. Complete any customization to the YAML file before you run the `kubectl create` command. By default the YAML files are generated in: `ibm-cloud-appmgmt-prod/ibm_cloud_pak/yaml/`

9. Create the persistent volumes for local storage by running the following command:

```

kubectl create -f install_dir/ibm-cloud-appmgmt-prod/ibm_cloud_pak/yaml/

```

10. Run the `pre-install.sh` script with the following parameters, note the following important choices first:

- Communication defaults to HTTP. Use the **--https** parameter to install either IBM Cloud App Management, Base or IBM Cloud App Management, Advanced with HTTPS communication. Changing between HTTP and HTTPS communication is not supported after installation.
- Use the **--advanced** parameter to install IBM Cloud App Management, Advanced. When you add the **--advanced** parameter HTTPS communication is enabled.
- To ensure a secure environment, you are encouraged to provide a valid **--cassandraUsername** other than the default "cassandra". When the **--cassandraUsername** parameter is passed to `pre-install.sh`, the script prompts you to enter a unique secure password to use for the Cassandra cluster rather than accepting the default value.

```

Usage ./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          Accept 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 <IP|FQDN> IP address or fully qualified domain name (FQDN)
for the ICP Master. You must use the address that you used to login to the ICP console.
In a highly available environment, this would be

```

```

the FQDN of the HAProxy or load balancer for ICP.
--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
--cassandraUsername <string> The username Cassandra will use. If left unset,
the default cassandra credentials will be used.

*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)
--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)

*Optional - Other
--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)
--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)
--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)

```

**Note:** The `--cassandraUsername`, can be set. If left unset, the default Cassandra credentials are used (default username "cassandra").

Example scenarios:

Scenario 1: No High Availability

```

ibm_cloud_pak/pak_extensions/pre-install.sh --accept --releasename ibmcloudappgmt --
namespace default --masterAddress icp-master.mydomain.com --proxyIP 0.1.2.3
--proxyFQDN icp-proxy.mydomain.com --clustercadomain mycluster.icp --advanced --
cassandraclustersize 1 --kafkaclustersize 1 --zookeeperclustersize 1 --couchdbclustersize 1
--datalayerclustersize 1 --elasticsearchclustersize 1 --minreplicashpas 1 --
maxreplicashpas 3 --cassandraUsername myCassandraUser

```

Scenario 2: High Availability with 3 Cassandra and 2 of each deployment (HPA minimums)

```

ibm_cloud_pak/pak_extensions/pre-install.sh --accept --releasename ibmcloudappgmt --
namespace default --masterAddress icp-master.mydomain.com --proxyIP 0.1.2.3
--proxyFQDN icp-proxy.mydomain.com --clustercadomain mycluster.icp --advanced --
cassandraclustersize 3 --kafkaclustersize 3 --zookeeperclustersize 3 --couchdbclustersize 3

```

```
--datalayerclustersize 3 --elasticsearchclustersize 3 --minreplicashpas 2 --
maxreplicashpas 3 --cassandraUsername myCassandraUser
```

### Scenario 3: High Availability with 6 Cassandra and 2 of each deployment (HPA minimums)

```
ibm_cloud_pak/pak_extensions/pre-install.sh --accept --releasename ibmcloudappmgmt --
namespace default --masterAddress icp-master.mydomain.com --proxyIP 0.1.2.3
--proxyFQDN icp-proxy.mydomain.com --clustercadomain mycluster.icp --advanced --
cassandraclustersize 6 --kafkaclustersize 3 --zookeeperclustersize 3 --couchdbclustersize 3
--datalayerclustersize 3 --elasticsearchclustersize 3 --minreplicashpas 2 --
maxreplicashpas 3 --cassandraUsername myCassandraUser
```

11. Optional: If you want to use your own certificate, enter the secret name manually when prompted. For more information, see [“Configuring a custom server certificate”](#) on page 104.
12. Optional: Review and further configure the Helm chart. For more information, see [“Configuring the Helm charts”](#) on page 119.
13. 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
```

14. Optional: If you want to either change the raw metric retention period from the default 8 days or enable metric summarization, add one (or both) of the following `--set` forms to the **helm install** command in step [“15”](#) on page 118:

```
--set global.metric.retention.rawMaxDays=2
```

where 2 represents the number of days to retain and can be a whole number from 2 to 32. Any value over 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 607.

This 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-1.5.0.tgz --tls`.

15. To deploy the Cloud App Management server Helm chart by using the Helm CLI, run the following command:

```
helm install --name my_release_name --values my_release_name.values.yaml
my_install_dir/ibm-cloud-appmgmt-prod-1.5.0.tgz --tls
```

Where `my_install_dir` is the directory where you extracted the Helm chart.

16. Run the `post-install-setup.sh` script with the following parameters, note the following important choices first:

- **--releaseName --instanceName --namespace**
- Use: **--advanced** for Advanced offering.

```
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"
--releaseName ${default_release} --instanceName ${instance_name} --namespace ${namespace}"
"example: for Advanced offering"
--releaseName ${default_release} --instanceName ${instance_name} --namespace ${namespace} --
advanced"
```



Example:

```
my_install_dir/ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions>/post-install-setup.sh --  
releaseName ibmcloudappgmt  
  
--namespace icam --instanceName ibmcloudappgmt --advanced
```

## Results

The Cloud App Management server is successfully installed.

## What to do next

1. Create a service instance. For more information, see [“Creating your service instance” on page 123](#).
2. Start the service instance and access the Cloud App Management console. For more information, see [“Starting the Cloud App Management UI” on page 124](#).

## 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 [“Installing IBM Cloud App Management on IBM Cloud Private” on page 99](#).

**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

Configuration setting	Value/Description/Example
	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”</a> on page 58.
<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> .
<code>global.persistence.storageClassOption.cassandra</code>	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 <code>default</code> .
<code>global.persistence.storageClassOption.zookeeper</code>	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 <code>default</code> .
<code>global.persistence.storageClassOption.kafka</code>	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 <code>default</code> .
<code>global.persistence.storageClassOption.datalayer</code>	Storage class for the Datalayer data. If you are using an environment-wide <code>global.persistence.storageClassName</code> ,

Configuration setting	Value/Description/Example
	such as <code>local-storage-datalayer</code> , leave this setting as default.
<code>global.persistence.storageClassOption.couchdbdata</code>	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.
<code>global.persistence.storageSize.cassandradata</code>	Storage size for the Cassandra data.
<code>global.persistence.storageSize.zookeeperdata</code>	Storage size for the ZooKeeper data.
<code>global.persistence.storageSize.kafkadata</code>	Storage size for the Kafka data.
<code>global.persistence.storageSize.datalayerdata</code>	Storage size for the Datalayer data
<code>global.persistence.storageSize.couchdbdata</code>	Storage size for the CouchDB data.
<code>ibm-cem.icpbroker:adminusername</code>	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 ':' '{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-1.5.0.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 `ibmcloudappgmt`. 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-1.5.0.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
ibmcloudappmgmt	default	Deployed	ibm-cloud-appmgmt-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
ibmcloudappmgmt-cassandra-0	1/1	Running	0	1d
ibmcloudappmgmt-couchdb-0	1/1	Running	0	1d
ibmcloudappmgmt-ibm-cem-datalayer-0	1/1	Running	0	1d
ibmcloudappmgmt-kafka-0	2/2	Running	1	1d
ibmcloudappmgmt-kafka-1	2/2	Running	0	1d
ibmcloudappmgmt-kafka-2	2/2	Running	0	1d
ibmcloudappmgmt-mongodb-0	1/1	Running	0	1d
ibmcloudappmgmt-zookeeper-0	1/1	Running	0	1d
ibmcloudappmgmt-zookeeper-1	1/1	Running	0	1d
ibmcloudappmgmt-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 on IBM Cloud Private”](#) on page 112.

### 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 *ibmcloudappmgmt*.

- `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” on page 14](#).

**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 on IBM Cloud Private” on page 112](#). 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 `kubect1 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 124](#).

## 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 123](#). 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 **ibmcloudappmgmt**.

## 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 16](#).

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 570](#).

You can start deploying your agents and data collectors as described in [“Step 6: Deploy the agents and data collectors” on page 52](#).

## 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 the Cloud App Management server

If you no longer want the Cloud App Management server on your system, you can delete the helm deployment.

## Procedure

Complete the following steps to uninstall the Cloud App Management server:

1. As the IBM Cloud Private administrator, log in to the management console and select your namespace. Run the following command:

```
cloudctl login -a my_cluster_URL --skip-ssl-validation
```

Where *my\_cluster\_URL* is the IBM Cloud Private name you defined for your cluster such as `https://<cluster_address>:443`. For any 'masterIP' reference, use the *cluster\_address* value.

**Note:** Run the following command to get your *cluster\_address* and port:

```
kubectl get configmap -n kube-public -o yaml
```

2. Find the Helm chart that you want to uninstall:

```
helm list --tls | grep ibm-cloud-appmgmt
```

3. Remove the Helm chart by running the following command:

```
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.



4. Delete the storage classes and persistent volume storage claims (PVCs) to release the claims on the persistent data store by issuing the following command:

```
kubectl delete storageclass --selector release=my_release_name
kubectl delete pvc --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.

5. Delete secrets and the cluster image policy by running the following command:

```
kubectl delete secrets --selector release=my_release_name --namespace my_namespace
kubectl delete clusterimagepolicy --selector release=my_release_name --namespace my_namespace
```

6. Optional: Back up the data on the persistent storage directories that you created on the worker nodes.

7. Delete the persistent volumes:

```
kubectl delete pv --selector release=my_release_name
```

8. Optional: You can safely remove the data from the persistent storage directories that you created on the worker nodes.

- /k8s/data/cassandra - Cassandra persistent storage

```
rm -r /k8s/data/cassandra
```

- /k8s/data/zookeeper - Zookeeper persistent storage

```
rm -r /k8s/data/zookeeper
```

- /k8s/data/kafka - Kafka persistent storage

```
rm -r /k8s/data/kafka
```

- /k8s/data/couchdb - CouchDB persistent storage

```
rm -r /k8s/data/couchdb
```

- /k8s/data/datalayer - Datalayer persistent storage

```
rm -r /k8s/data/datalayer
```

- /k8s/data/elasticsearch - Elasticsearch persistent storage

```
rm -r /k8s/data/elasticsearch
```

9. Optional: You can remove the Cloud App Management image from IBM Cloud Private. For more information, see the [Removing an image from the console](#) topic in the IBM Cloud Private Knowledge Center.

## Results

The Cloud App Management server helm chart is uninstalled. The storage configuration that was required for the installation was also deleted.

## 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 128.

For instructions on backing up and restoring CouchDB, see [“Back up and restore CouchDB”](#) on page 129.

For instructions on backing up and restoring ZooKeeper, see [“Backup restore ZooKeeper”](#) on page 130.

## Back up and restore Cassandra

The following procedure describes how to back up 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 “8” on page 114 in the *Installing the Cloud App Management server* topic. 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.

### 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 backup.

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 issuing the command below:

```
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`. These are the directories you prepared in the [“Installing on IBM Cloud Private”](#) on page 112 topic.

You will 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. We recommend that you 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
```

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

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 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 2019.2.1 is to specify keyspaces individually.

Verify the restore is successful

5. Ensure there are no failure messages in the restore logs, and that the Cloud APM console displays 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.

## Backup restore ZooKeeper

Use the following procedure to back up ZooKeeper.

### About this task

#### Procedure

Back up ZooKeeper

1. Determine the IP address of the config service. If you restart the config service, repeat this step before you back up.

```
kubectl get service --all-namespaces |grep ibmcloudappmgmt-config  
#List the service for ibmcloudappmgmt-config
```

You see a result similar to:

```
ibmcloudappmgmt-config ClusterIP 10.0.0.133 <none> 80/TCP,443/TCP 6h33m
```

2. Back up the ZooKeeper data by running the following command on the master node:

```
curl http://icam_config_service_ip/1.0/systemconfig/backup >icam_config_backup
```

Where

*icam\_config\_service\_ip* is the config service IP address, for example, 10.0.0.133  
*icam\_config\_backup* is the JSON file name that you want to use for backup, for example config-data.json

When the backup completes, you see a result similar to:

% Total	% Received	% Xferd	Average Speed	Time	Time	Time	Current				
			Dload	Upload	Total	Spent	Left	Speed			
100	6539	100	6539	0	0	3069	0	0:00:02	0:00:02	--:--:--	3069

The output file contains data in JSON format, for example: {"name" : "value"}

Restore Zookeeper

3. Restore the ZooKeeper data

```
curl -i -X POST --header "Content-Type: application/json" -d @icam_config_backup "http://icam_config_service_ip/1.0/systemconfig/backup"
```

Where

*icam\_config\_backup* is the ICAM config backup JSON file name, for example config-data.json  
*icam\_config\_service\_ip* is the ICAM config service IP address, for example, 10.0.0.133

Verify that the ZooKeeper data is restored

4. Run the following command to verify that the ZooKeeper data is restored:

```
curl -X GET http://10.0.0.133:80/1.0/systemconfig/backup >config-data-afterRestore.json
```

5. To validate if your restore was successful, perform a diff between config-data.json (this is the file name that you specified for *icam\_config\_backup* in step "2" on page 130 and config-data-afterRestore.json. The content should be identical.

---

## Chapter 9. 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 13, “Integrating with other products,”](#) on page 517.

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 132.
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 132.

**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 134
  - [“Installing agents on Linux systems”](#) on page 138
  - [“Installing agents on Windows systems”](#) on page 141

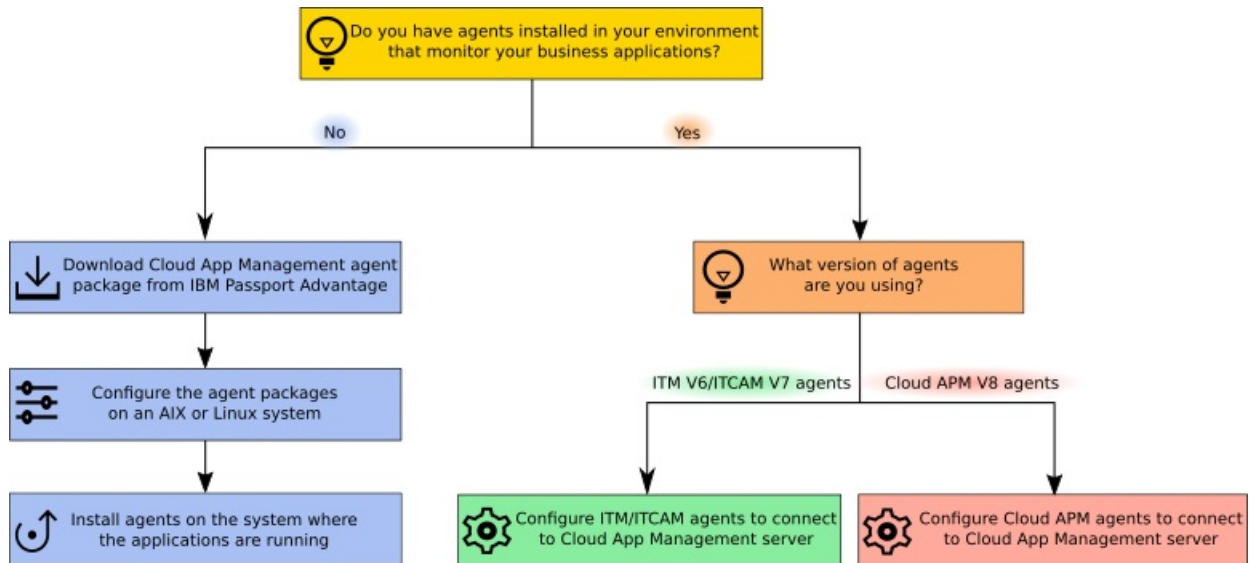
**Remember:** If you install the agent as a non-root user on the AIX or Linux system, run the `UpdateAutoRun.sh` script with root user or sudo user access after a non-root installation. See [“Installing agents as a non-root user”](#) on page 145.

---

## 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 132](#)
2. [“Configuring the downloaded images” on page 132](#)
3. [Installing the agents](#)
4. [“Integrating with IBM Tivoli Monitoring agents” on page 517](#)
5. [“Integrating with Cloud APM, Private agents” on page 530](#)

## 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 53](#).
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 132](#).
- For the Data Collectors for Cloud Resources see [Chapter 11, “Deploying ICAM Data Collectors,” on page 427](#).

## 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 on IBM Cloud Private” on page 112](#). Also, familiarise yourself with the security certificates that are available, see [“Configuring certificates for HTTPS communications” on page 104](#).

## 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 `preconfig.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 134](#)
- [“Installing agents on Linux systems” on page 138](#)
- [“Installing agents on Windows systems” on page 141](#)

## 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:

- 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
- 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
- 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 57](#) 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 145](#).



### 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 57 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 145.

#### 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 134](#).

### Before you begin

- Review the information in [“System requirements” on page 57](#) 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 132](#).
- Review the agent preinstallation tasks before you install the agents. See [“Preinstallation on AIX systems” on page 134](#).
- Configure the agent images with the connection details for the Cloud App Management server. See [“Configuring the downloaded images” on page 132](#).

**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 145](#). 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 144](#).

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 132](#).

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 145](#).
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 146](#) or if you want to review the default settings, see [Chapter 10, “Configuring the ICAM Agents,” on page 161](#).
- 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 162.

- 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 124.

## 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:

- Cisco UCS agent
- DataPower agent
- Db2 agent
- DataStage agent
- Hadoop agent
- HTTP Server agent
- IBM Integration Bus agent
- JBoss agent
- Linux OS agent
- Linux KVM agent
- MongoDB agent
- MySQL agent
- Oracle Database agent
- PostgreSQL agent
- SAP agent
- SAP HANA Database agent
- SAP NetWeaver Java Stack agent
- Tomcat agent
- VMware VI 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 57 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 145.

## 70-character limitation for installation path

The installation directory and the path to it must be no more than 70 characters.

## 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 149](#).

**Note:** You do not need to do this step for data collectors.

## 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.

### 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**.

## 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 138](#).

### Before you begin

- Review the information in [“System requirements” on page 57](#) 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 132](#).
- Review the agent preinstallation tasks before you install the agents. See [“Preinstallation on Linux systems” on page 138](#).

- Configure the agent images with the connection details for the Cloud App Management server. See [“Configuring the downloaded images” on page 132](#).

**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 145](#). 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 144](#).

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 132](#).

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 145](#)).

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 146 or if you want to review the default settings, see Chapter 10, [“Configuring the ICAM Agents,”](#) on page 161.
- 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 162.

- 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 124.

## 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.

- Cisco UCS agent
- Db2 agent
- DataStage agent
- Hadoop agent
- HTTP Server agent
- IBM Integration Bus agent
- JBoss 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
- SAP agent
- SAP HANA Database agent
- SAP NetWeaver Java Stack agent

- Skype for Business Server agent
- Tomcat agent
- VMware VI 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 57 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 142.

### Before you begin

- Review the information in [“System requirements”](#) on page 57 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 132.
- Review the agent prerequisite tasks before you install the agents. For details, see [“Preinstallation on Windows systems”](#) on page 142.
- Configure the agent images with the connection details for the Cloud App Management server. See [“Configuring the downloaded images”](#) on page 132.

### 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\_2019.3.0.

```
cd extracted_image_directory
installAPMAgents.bat
```

To install the agents in silent mode, see [“Installing agents silently” on page 144](#).

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 132](#).

**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 146 or if you want to review the default settings, see Chapter 10, [“Configuring the ICAM Agents,”](#) on page 161.
- 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 162.

- 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 124.

## 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 138 and [“Installing agents on UNIX systems”](#) on page 134.
2. Optional: If you installed your agent as a selected user and want to configure the agent as a different user, run the `./secure.sh` script.

For example:

```
./secure.sh -g mqadmin1
```

For more information about the `./secure.sh` script, see [“Configuring agents as a non-root user”](#) on page 167 and [Securing the agent installation files](#).

3. Optional: Configure your monitoring agents on Linux or AIX as necessary, see [Chapter 10, “Configuring the ICAM Agents,”](#) on page 161.
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 167.

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 166.

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 168.

## 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 11](#) on page 146.

For information about how to configure an agent, see [Chapter 10, “Configuring the ICAM Agents,”](#) on page 161.

Agent	Configured and started automatically	Configured manually and started automatically	Configured and started manually	Multiple instance (started manually)
Cisco UCS agent	—	✓	✓	✓
DataPower agent	—	—	—	✓
DataStage agent	—	—	✓	✓
Db2 agent	—	—	—	✓

Table 11. Agent postinstallation checklist (continued)

Agent	Configured and started automatically	Configured manually and started automatically	Configured and started manually	Multiple instance (started manually)
IBM Integration Bus agent	—	—	—	✓
	—	—	—	✓
Linux OS agent	✓	—	—	—
Linux KVM agent	—	—	✓	✓
Microsoft Cluster Server agent	—	—	✓	—
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	—	—	—	✓
SAP agent	—	—	✓	✓
SAP HANA Database agent	—	—	✓	✓
SAP NetWeaver Java Stack agent	—	—	✓	✓
Skype for Business Server agent	—	✓	✓	✓
Tomcat agent	—	—	✓	✓
UNIX OS agent	✓	—	—	—
VMware VI agent	—	—	✓	✓

Table 11. 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 162](#). For a list of multiple-instance agents, see [Table 11 on page 146](#).



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 152.

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 the following command 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 162.

- 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` 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 153. 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 416. 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 152 or [“Unconfiguring the data collector in silent mode”](#) on page 153.

#### 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 155.
- To manually unconfigure the data collector for the Liberty server, see [“Manually unconfiguring the data collector for WebSphere Application Server Liberty”](#) on page 156.

#### *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:

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 419](#).

#### 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.sh/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 10. 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 162](#).

### 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 13 on page 162](#).
- *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 13 on page 162](#).
- `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 13 on page 162](#).

Monitoring agent	name	Two letter agent code
Cisco UCS agent	ciscoucs	v6
DataPower agent	datapower	bn
Db2 agent	db2	ud
DataStage agent	datastage	td
Hadoop agent	hadoop	h8
Microsoft Hyper-V Server agent	hyper-v	hv
IBM Integration Bus agent	iib	qi
JBoss agent	jboss	je
Linux OS agent	os	lz
Linux KVM agent	linux_kvm	v1
Microsoft .NET agent	dotnet	qe

Table 13. Agent names and agent codes (continued)

Monitoring agent	name	Two letter agent code
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
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	q1
Tomcat agent	tomcat	ot
UNIX OS agent	os	ux
VMware VI agent	vmware_vi	vm
WebSphere Applications agent	was	yn
WebSphere Infrastructure Manager agent	wim	d0
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 14 on page 164](#), where *name* is the agent name that is specified in [Table 13 on page 162](#).

Table 14. 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 <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 11 on page 146</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.
<code>./name-agent.sh uninstall</code>	Uninstalls the monitoring agent. For more information, see <a href="#">“Uninstalling your agents” on page 150</a> .
<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 15 on page 165](#), where `name` is the agent name that is specified in [Table 13 on page 162](#).

Table 15. 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 11 on page 146</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.
<code>name-agent.bat uninstall</code>	Uninstalls the monitoring agent. For more information, see <a href="#">“Uninstalling your agents” on page 150</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

```

```
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 136](#) on AIX systems or [“Installing agents” on page 139](#) on Linux systems.
2. Optional: Configure your monitoring agents on Linux or AIX as necessary, see [Chapter 10, “Configuring the ICAM Agents,” on page 161](#).
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/kcirunas.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 13 on page 162](#). 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 138](#) and [“Installing agents on UNIX systems” on page 134](#).
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 10, “Configuring the ICAM Agents,” on page 161](#).
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 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 57](#) 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 [Creating a user and granting required permissions](#).
- 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 173](#).

### 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 172](#)
  - [“Configuration parameters for the data provider” on page 172](#)
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 124.
- 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 174.

## 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 172

– [“Configuration parameters for the data provider”](#) on page 172

- 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 124](#).
- 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 174](#).

## 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 172](#)
- [“Configuration parameters for the data provider” on page 172](#)

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 124](#).

- 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 174.

## 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.

Parameter name	Description	Mandatory field
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
Validate SSL Certificates	A Boolean value that indicates whether the agent validates SSL certificates when the agent uses SSL to communicate over the network. 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. <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 173.	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.

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

Table 17. Name and description of the configuration parameters for the data provider (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	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/1x8266/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 `$CANDLEHOME/lx8266/vm/bin/kv6_data_provider.sh` file.
2. Add the following line before the line that starts with `KV6_JVM_ARGS="$KV6_CUSTOM_JVM_ARGS...`:

```
KV6_CUSTOM_JVM_ARGS=-Xmx1024m
```

3. Restart the agent.

## Configuring 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 37](#).



## 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](#).

To see monitoring data, such as resource utilization, throughput, and connection statistics, enable resource monitoring. For instructions, see [“Resource monitoring”](#) on page 175.

**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. The default IP address is *9.123.109.139*.

#### 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 176 and “5” on page 177 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 166](#).

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 611](#).

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 184](#).

- 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 124.

## 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 187.

### 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 184](#).
- 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 124](#).

## 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 187](#).

### 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<sup>®</sup> 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/sqlllib/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 184](#).
- 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 124](#).

## 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 187](#).

### 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 184](#).
- 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 124](#).

## 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 185.

### 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\_DUMPOUT** 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 53](#).

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, ResourceManager, 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 635](#).

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 196](#).

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 37](#).

## 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 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**.

6. 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**.
7. 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**.
8. 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**.
9. 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.
10. Click **OK**. The specified configuration settings are saved.
11. 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 197](#).
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 124](#).

## 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 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.

5. 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.
6. 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):
7. 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.
8. 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)
9. 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:)
10. 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 ResourceManger 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 ResourceManger that you want to remove.

11. 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.

12. 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 197.
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 124.

## 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) 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)
```

- c) Optional: For the **HADOOP\_ADDITIONAL\_NN\_HOSTNAME** parameter, specify the host name of the Standby or Secondary NameNode.
- d) 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.

- e) Optional: For the **Hadoop\_SRM** parameter, type Yes to add Standby ResourceManagers for a high-availability cluster, and go to [step g](#).
- f) 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
```

- g) For the **JAVA\_TRACE\_LEVEL** parameter, specify the appropriate trace level.
  - h) Optional: For the **JAVA\_JVM\_ARGS** parameter, specify arguments for the Java™ virtual machine.
  - i) 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 197.
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 124.

## 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 37](#).

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\temp\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\temp\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 37](#).
- 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 MQ 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 201](#).
  - 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 204](#).

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 208.

## 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 201.
- To configure the agent without interaction, edit the silent response file and then run the configuration script. See [“Silent configuration”](#) on page 202.

### 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 V10, specify the installation directory of IBM Integration Bus V10 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 124](#).

**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 204](#).
- To enable resource statistics data for an integration node (broker), see [“Enabling JVM resource statistics” on page 207](#).

### 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 208.](#)

## 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** | **UNIX** `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 37.

## 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 211.

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 211.

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\KDTEENV_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 211](#).

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 124](#).

## 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.

<i>Table 19. Names and descriptions of the metadata repository configuration parameters</i>		
<b>Parameter name</b>	<b>Description</b>	<b>Mandatory field</b>
Database Type	Database type of the metadata repository. The Db2 default value is 1.	Yes
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 20. 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 21. 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 214](#).
  2. [“Add a JBoss Server Management User” on page 213](#).
  3. [“Enabling Web/HTTP Statistic Collection” on page 215](#). 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 218](#).
  - [“Configuring the agent by using the silent response file” on page 219](#).
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 218](#).
  - [“Configuring the agent by using the silent response file” on page 219](#).

### 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 124](#).

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>
  ...
</interfaces>
```

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>
  ...
</interfaces>
```

## 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 23 on page 217](#) - Example JMX service URLs.

Parameter name	Description	Silent configuration file parameter name
Server Name	Provide a name to identify the JBoss/WildFly Server.	<b>KQE_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 23 on page 217</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:NNNN}" />. 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 22 on page 217](#) for an explanation of each of the configuration parameters.
5. Enter the Java settings, then click **Next**.  
See [Table 22 on page 217](#) for an explanation of each of the configuration parameters.
6. Enter the JMX settings, then click **Next**.  
See [Table 22 on page 217](#) 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 22 on page 217](#) 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 217 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:
  - a) 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
```

- b) 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 217 for an explanation of each of the configuration parameters.

- c) 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 227](#)
- [“Configuring a connection to the RHEVM server” on page 225](#)

## 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 222, “4” on page 222, and “5” on page 222 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 `LIBVIRTD_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 227.

## 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 221.

**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 227.

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 124.

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 221.
- 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 221.

### 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 229.

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 124.

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.

Parameter name	Description	Mandatory field
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 24. 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 24. Names and descriptions of the configuration parameters for connecting to the RHEVM server (continued)

Parameter name	Description	Mandatory field
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 25. Names and descriptions of the configuration parameters for connecting to the hypervisor

Parameter name	Description	Mandatory field
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 25. 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 25. 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 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 37.

### 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 following groups on 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 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 .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 124.

## 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 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 37.

### 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 234.

## 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 124](#).

## 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 37](#).
- Ensure that you complete the following tasks:
  - [“Creating users” on page 235](#)
  - [“Assigning administrator rights to the Exchange Server user” on page 238](#)
  - [“Making the Exchange Server user a local administrator” on page 238](#)
  - [“Configuring the Exchange Server for reachability” on page 240](#)
  - [“Configuring the agent to run under the domain user” on page 241](#)

### 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 241](#).
- To configure the agent by using the silent response file, see [“Configuring the agent by using the silent response file” on page 245](#).

## 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 236](#)
- [“Creating users by running the New User utility” on page 237](#)

## 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 238](#)
- [“Assigning administrator rights on Exchange Server 2010” on page 238](#)

### 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 239](#)
- [“Making the user a local administrator on Windows 2008 computer” on page 239](#)
- [“Making the user a local administrator on Windows 2012 computer” on page 239](#)
- [“Making the user a local administrator on Windows 2016 computer” on page 240](#)

### **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 242](#)
- [“Configuration parameters for Exchange services” on page 243](#)
- [“Configuration parameters for reachability” on page 244](#)

For information about the validation of configuration values, see [“Validation of configuration values” on page 245](#).

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 124](#).

### 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.

<i>Table 26. Names and descriptions of configuration settings in the Exchange Server Properties tab</i>			
<b>Parameter name</b>	<b>Description</b>	<b>Mandatory field</b>	<b>Examples</b>
Exchange Server Name	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.  <b>Remember:</b> In clustered and distributed environments, specify the Mailbox Server name for Exchange Server 2007.	Yes  <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.	If the Exchange Server name is popcorn, enter popcorn in the <b>Exchange Server Name</b> field.
Exchange Domain Name	The name of the domain where the Exchange Server is installed.	Yes	If the Exchange Server is in the LAB.XYZ.com domain, enter the name that precedes the first dot, for example, LAB.
Exchange User Name	The name of the user who is configured to access the Exchange Server.  <b>Remember:</b> The user must have a mailbox on the same Exchange Server.	Yes	

Table 26. Names and descriptions of configuration settings in the Exchange Server Properties tab (continued)

Parameter name	Description	Mandatory field	Examples
Exchange User Password	The password of the user who is configured to access the Exchange Server.	Yes	
Confirm Password	The same password that you specified for the Exchange Server user.	Yes	
Exchange MAPI Profile Name	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.	No	
Configuration in cluster	Select this check box if you want to configure the Microsoft Exchange Server agent in a cluster environment.	Not applicable	
Cluster Server Name	The name of the Cluster Server. This field is enabled when you select the <b>Configuration in cluster</b> check box.	Yes, if the field is enabled.	SCCCluster
Exchange Subsystem ID	The name of the Cluster Server node. This field is enabled when you select the <b>Configuration in cluster</b> check box.	Yes, if the field is enabled.	node1
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 27. 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

Table 27. Names and descriptions of configuration settings in the Exchange Services Monitoring tab (continued)

Parameter name	Description	Mandatory field
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 28. 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.
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

Table 28. Names and descriptions of configuration settings in the Advanced Configuration Properties tab (continued)

Parameter name	Description	Mandatory field
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 29. Validation tests	
Validation test	Verifies whether
Exchange Server Name	The Mailbox Server name of the user matches the specified Exchange Server name.
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 124.

## 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 124](#).

## 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 248.
  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 249.
  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 250.
  7. To modify the DCOM security permission for a non-administrator user, refer [“Modifying DCOM permissions”](#) on page 249.
  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 37](#).
- 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 252](#).

## 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 124](#).

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 124.

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 162. For detailed information about the agent version list and what's new for each version, see [“Change history”](#) on page 37.

- 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 253.

- 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 form 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 256](#).
- Change the user account from the local user to the domain user. For more information, see [“Changing the user account” on page 256](#).
- 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 124](#).

## 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 30.*

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 256](#).
- Change the user account from the local user to the domain user. For more information, see [“Changing the user account” on page 256](#).
- 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 124](#).

## 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 = alam@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 258.

## 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:

<b>Environment variables</b>	<b>Description</b>
<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.
<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 57](#) 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 259](#).

## 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 124.

## 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 261.
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 37](#).

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 **sysdamin** 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 262](#)

Use one of the following procedures to create a user:

- [“Creating an SQL Server user ID with Windows authentication” on page 263](#)
- [“Creating an SQL Server user ID with SQL Server authentication” on page 264](#)

Use the following procedure to grant permissions:

- [“Granting minimum permissions for data collection” on page 264](#)
- [“Granting permission to the Perflib registry key for collecting data for few data sets” on page 266](#)

### Checking the permissions of an existing SQL Server user

You can run the utility tool **koqVerifyPerminssions.exe** to check if an existing SQL Server user has sufficient permissions related to SQL Server databases.

### About this task

The utility tool **koqVerifyPerminssions.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* \TMAITM6where *Agent\_home* is the agent installation directory.
2. Run the **koqVerifyPerminssions.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 264.

**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 264.

**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 32. 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.

Table 32. Parameters (continued)	
Parameters	Description
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 268](#).
  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 270.

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 Parm**s 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 124.

### 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.



Table 33. Names and descriptions of 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 33. 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 33. 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  <math>\\$instance\_name</math>  \LOG\ERRORLOG,  where <i>instance_name</i> is the Microsoft SQL Server instance name.</p>

Table 33. 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 33. 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 33. 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 33. 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 33. 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 279](#).



## 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.ctl file. Complete the following steps to modify the default limit:

1. Specify the maximum number of rows for KOQFGRPD in the koqcoll.ctl 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.ctl file.

If the value in the koqcoll.ctl file is less than the value that is specified in the **COLL\_KOQFGRPD\_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\_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%\_%\_DEVD\_TEMP file to the %COLL\_HOME%\_tmp\_%COLL\_VERSION%\_%COLL\_SERVERID%\_%COLL\_SERVERID%\_%\_DEVD\_PREV file.

%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\_WRAP\_LINES:** 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 124](#).

## 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 278](#).

### 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 286.

### 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 286.

### 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. Refer the following prerequisites for configuring MongoDB agent for both remote and local 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 34 on page 291](#).

- 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:

Table 34. Attributes groups and their required user roles

Roles	MongoDB database version	Attribute groups
dbAdminAnyDatabase	2.x, 3.x, 4.x	Response Times
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 162](#). For detailed information about the agent version list and what's new for each version, see the [“Change history” on page 37](#).

### 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 124](#).

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 34 on page 291](#).
  - 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 124](#).

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 34 on page 291](#).
  - 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 124](#).

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 37](#).

### 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 124.

## 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 296.

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 124.

## 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 124.

## 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 298.

### 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/tmaintm6`

- 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 301](#)
  - [“Configuration parameters for the OnCommand Unified Manager” on page 302](#)
  - [“Configuration parameters for the OnCommand API Service” on page 303](#)
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 301](#)
- [“Configuration parameters for the OnCommand Unified Manager” on page 302](#)
- [“Configuration parameters for the OnCommand API Service” on page 303](#)

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 301](#)
- [“Configuration parameters for the OnCommand Unified Manager” on page 302](#)
- [“Configuration parameters for the OnCommand API Service” on page 303](#)

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.

Parameter name	Description	Mandatory field
Instance Name ( <b>KNU_INSTANCE_NAME</b> )	The name of the instance. <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 ( <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 ( <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 35. Name and description of the configuration parameters for the data provider (continued)

Parameter name	Description	Mandatory field
Level of Detail in Data Provider Log ( <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 36. Name and description of the configuration parameters for the OnCommand Unified Manager

Parameter name	Description	Mandatory field
Server ( <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 ( <b>KNU_DATASOURCE_PASSWORD</b> )	The password of the user that you specify in the <b>User</b> parameter.	Yes



Table 36. 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 ( <b>KNU_DATASOURCE_</b> <b>PROTOCOL</b> )	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 37. Name and description of configuration parameters for the OnCommand API Service

Parameter name	Description	Mandatory field
Host Address ( <b>KNU_API_SERVICES_HO</b> <b>ST_</b> <b>ADDRESS</b> )	The host name or IP address of the OnCommand API service.	Yes
User ( <b>KNU_API_SERVICES_</b> <b>USERNAME</b> )	A user name with sufficient privileges to connect to the OnCommand API service. The default value is admin.	Yes
Password ( <b>KNU_API_SERVICES_</b> <b>PASSWORD</b> )	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 305.](#)
  - [“Configuring the agent by using the silent response file” on page 312.](#)
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 308.](#)
  - [“Configuring the agent by using the silent response file” on page 312.](#)

## 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 124](#).

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 124](#).

## 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 310.

- 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 310, 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 124](#).

## 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 124.](#)

## 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 315](#) and [“Running the krzgrant.sql script” on page 315.](#)

**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 317](#).

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 166](#).

#### 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 321 and “Format file ” on page 330. You must ensure that you save these files on the system where you access the Performance Management console so that you can upload the files to the Cloud APM server.

#### About this task

To enable the OS agents to monitor log files, you must upload the configuration file and format file and specify to which OS agent the configuration applies. The OS agent downloads the `.conf` and `.fmt` files and the agent monitors the log files that you specify in the configuration.

#### 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 333.

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 18, “Event filtering and summarization,”](#) on page 665.

### 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 18, "Event filtering and summarization," on page 665](#).

## **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 18, "Event filtering and summarization," on page 665](#).

## **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 330.

### **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 338.

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 338.

**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 19, “Windows Event Log,”](#) on page 667.

### 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 19, “Windows Event Log,”](#) on page 667.

### WINEVENTLOGS

Controls which Windows event logs are monitored.

The `WINEVENTLOGS` statement is a comma-delimited list with no spaces. For more information, see [Chapter 19, “Windows Event Log,”](#) on page 667.

**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 333.

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 321](#).

**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 330](#).

### 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 335](#).

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 18, “Event filtering and summarization,” on page 665](#). `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 321](#).

### 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 37.

## 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 124](#).

## 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 124](#).

## 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 124](#).



## 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 162](#). For detailed information about the agent version list and what's new for each version, see the [“Change history” on page 37](#).

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 345](#).

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 124](#).

## 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 345](#).

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 124](#).

## 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 345](#).

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 124](#).

## 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 37](#).
- 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 346](#)
  - [“Configuring the agent by using the silent response file” on page 348](#)
- 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 347](#)
  - [“Configuring the agent by using the silent response file” on page 348](#)

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 351](#).

After you configure the SAP agent, you must verify the agent configuration. For more information, see [“Verifying agent configuration” on page 358](#).

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 357](#).

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 348](#)

- 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 348](#)

**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 348](#).
  - 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 348](#).

**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 348](#).

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 40. 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 40. 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 40. 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. 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> =true or false
SNC Security Level	The security level of SNC.	Yes	<b>sap_snc_mode1.sap_snc_qop</b> =QOP value. 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> = 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> = Server SNC Name
SAP Cryptolibrary Path	The path of SAP Cryptolibrary.	Yes	<b>sap_snc_mode1.sap_snc_library</b> = 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 353.

### 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 41. 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 357](#).
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\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 42. 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 42. 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: **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 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**.
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 to import the delete transport:
  - 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.

## 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 358](#)
- [“Copying the NetWeaver RFC SDK V7.20 library in SAP agent setup” on page 359](#)
- [“Verifying the NetWeaver RFC SDK V7.20 library” on page 359](#)
- [“Verifying the configuration of Solution Manager V7.1 with MAI-Monitoring” on page 360](#)
- [“Verifying MAI Alerts” on page 361](#)
- [“Verifying configuration settings specific to third-party component” on page 361](#)

### 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
  - libicudcnumber.dll
  - libsapucum.dll
  - icudt34.dll
  - icuuc34.dll
  - sapnrwfc.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/intrp/sa/lib* location, where *intrp* is the operating system code (aix526, li6263, sol606). You must copy the following files:
  - libsapnrwfc.so
  - libicudcnumber.so
  - libicuuc34.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 `sapnrwfc.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 libsapnrwfc.so | grep SAPFileVersion`**
  - c) You must see the following message: `[root@IBMSAP2V6 lib]# strings libsapnrwfc.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 362](#)
- [“SAP user IDs” on page 363](#)
- Utilities for the SAP agent
- [“SAP RFC connections” on page 363](#)
- [“Test Connection feature” on page 373](#)
- [“Optional advanced configuration in SAP” on page 365](#)
- [“CEN CCMS reporting” on page 371](#)
- [“Uninstalling the Advanced Business Application Programming \(ABAP\) transport from the SAP system” on page 372](#)

**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 363](#)
- [“Basic agent monitoring” on page 363](#)

### **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 43. 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 43. 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 366](#)
- [“Copy, back up, and restore data by using transactions” on page 366](#)
- [“Command-line utility tool” on page 367](#)
- [“Running the command-line utility on a Windows environment” on page 368](#)
- [“Running the command-line utility on a Non-Windows environment” on page 368](#)
- [“Alerts maintenance” on page 369](#)
- [“Select monitor sets and monitors transaction” on page 369](#)
- [“Configure Dialog Step Response Threshold in the SAP system” on page 370](#)

**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 370

### ***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 368 and [“Running the command-line utility on a Non-Windows environment”](#) on page 368.

#### • 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 351.

**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.

<i>Table 44. /IBMMON/ITM_CNFG</i>	
<b>PARAM_NAME</b>	<b>VALUE_CHAR</b>
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.

<i>Table 45. /IBMMON/ITM_CNFG</i>	
<b>PARAM_NAME</b>	<b>VALUE_CHAR</b>
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 46. /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 47. /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 375](#)
- [“Configuring the agent on Linux and AIX systems” on page 376](#)
- [“Configuring the agent by using the silent response file” on page 376](#)

## 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 124](#).

## 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 124.

## 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 124.

## 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 37.

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 377 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 380.
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 124.

## 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 377 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 380](#)

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 124](#).

## 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 377](#) 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 380](#)

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 124](#).

## 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 has 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 information about how to check the version of an agent in your environment, see . For detailed information about the agent version list and what's new for each version, see the [“Change history” on page 37](#).

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 383](#).

## 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 384.

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 124.



## 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 384](#).

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 124](#).

## 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 49. 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.

<i>Table 49. Names and descriptions of the configuration parameters for the agent (continued)</i>	
<b>Parameter name</b>	<b>Description</b>
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  <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. 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.

## 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 37](#).

## 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 167](#). 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 166](#).

## 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 124](#).

## 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 124.

## 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 124.

## 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 124.

## 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 391.
- 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 390.

## 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> (2 GB)
<b>Medium environment</b>	A vCenter server that manages between 125 - 250 ESX(i) servers and 1500 - 4000 guests.	<b>-Xmx4096m</b> (4 GB)
<b>Large environment</b>	A vCenter server that manages between 250 - 500 ESX(i) servers and 4000 - 7500 guests.	<b>-Xmx8192m</b> (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> (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 396.

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 124.

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 396.

## 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:

- 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.

- Respond to the prompts by referring to the following topics:

- [“Configuration parameters for the data provider”](#) on page 395
- [“Configuration parameters for the data source”](#) on page 394

- 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 124.

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 396.

## 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:

- 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 395](#)
- [“Configuration parameters for the data source” on page 394](#)

- 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 124](#).

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 396](#).

## 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.

Table 51. Names and descriptions of the configuration parameters for the data source

Parameter name	Description	Mandatory field
Data Source ID	The ID of the data source.	Yes
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 52. 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

Table 52. Names and descriptions of the configuration parameters for the data provider (continued)

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 391.</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 390.

**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 37](#).

### 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 397](#) 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 400](#).
- (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 401](#).
- (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 407](#).
- (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 401](#).
- (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 416](#).

### 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 401.

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 124](#).

## 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 399](#).

### 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 124](#).

## 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 399](#).

**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 402.

- 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 405.

- For silent configuration, see [“Configuring the data collector in silent mode”](#) on page 407.

### 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 403.
  - For WebSphere Application Server Liberty, proceed to Step “10” on page 403.
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 412.
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 404. 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 **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 “20” on page 404. To use the log path that is already specified in the JVM argument of the application server, enter 2.
19. 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`.

20. 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.
21. Enter a to accept your changes.
22. 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.

23. 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 403.
  - 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 406. 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 406.
  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 407. 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 406.
  - 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 409.

## 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 405.
- 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 153.

### 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 53 on page 409](#).

#### **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 53 on page 409](#) 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 53. 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>	
ttapi.enable	Specifies whether the data collector supports transaction tracking. Valid values are True and False. <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.

Table 53. Available properties for running the configuration utility in silent mode (continued)

Property	Comment
<b>PMI resource and data collector monitoring</b>	
<p>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.</p>	
tema.appserver	<p>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.</p> <p>When this parameter is set to False, diagnostics and transaction tracking features are not available, and only resource monitoring data is collected.</p>
tema.jmxport	<p>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.</p>
<b>Integration of the data collector with the monitoring agent component of WebSphere Applications agent</b>	
temaconnect	<p>Specifies whether the data collector connects to the monitoring agent component of WebSphere Applications agent. Valid values are True and False.</p> <p><b>Important:</b> You must use the True value to use the WebSphere Applications agent.</p>
tema.appserver	<p>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.</p> <p>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.</p>
tema.host	<p>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).</p>
tema.port	<p>Specifies the port number of the monitoring agent component of WebSphere Applications agent. Do not change the default value of 63335.</p>
tema.jmxport	<p>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.</p>

Table 53. Available properties for running the configuration utility in silent mode (continued)

Property	Comment
<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	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.  <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> .
<b>WebSphere Administrative Services connection settings</b>	
was.wsadmin.connection.host	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.  <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.
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.
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.

Table 53. Available properties for running the configuration utility in silent mode (continued)

Property	Comment
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	<p>Specifies the application server instance within the application server profile to configure.</p> <p><b>Tip:</b></p> <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	<p>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.</p> <p><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.</p> <p><b>Tip:</b> The silent response file can have multiple instances of this property.</p>

### 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/KYENV`
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>

```

## 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 416](#).
- To manually configure the data collector for the Liberty server, see [“Manually configuring the data collector for WebSphere Application Server Liberty” on page 418](#).

### 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=${WAS_SERVER_NAME} -Xbootclasspath/p:${ITCAMDHOME}/  
toolkit/lib/bcm-bootstrap.jar -Djava.security.policy=${ITCAMDHOME}/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.

<i>Table 54. Environment entry</i>		
<b>Platform</b>	<b>Environment entry name</b>	<b>Environment entry value</b>
AIX R6.1 (64-bit JVM)	LIBPATH	/lib:\${ITCAMDHOME}/ toolkit/lib/aix536
AIX R7.1 (64 bit JVM)	LIBPATH	/lib:\${ITCAMDHOME}/ toolkit/lib/aix536
Linux Intel R2.6 (32-bit JVM)	LD_LIBRARY_PATH	/lib:\${ITCAMDHOME}/ toolkit/lib/li6263
Linux x86_64 R2.6 (64-bit JVM)	LD_LIBRARY_PATH	/lib:\${ITCAMDHOME}/ toolkit/lib/lx8266
Windows (32-bit JVM)	PATH	/lib;\${ITCAMDHOME}/ toolkit/lib/win32
Windows (64-bit JVM)	PATH	/lib;\${ITCAMDHOME}/ 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 *ITCAMDHOME* variable. Set the *ITCAMDHOME* 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 155.

### 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.

Table 55. 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

Table 55. Environment entry (continued)		
Platform	Environment entry name	Environment entry value
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

- 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>
```

- 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
```

- Change the am.camtoolkit.gpe.dc.operation.mode line as follows and save your changes.

```
am.camtoolkit.gpe.dc.operation.mode=WR
```

- 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 155.

## 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 157.

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 124](#).

## 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 37](#).
- 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 422](#).

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 423.
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 424.

## 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:
  - a) Log on to the Windows systems as a system administrator.
  - b) Add the user ID that is used to run the agent to the **mqm** group.
  - c) 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:
    - a. Click **Start** > **Run**, and then type `regedit.exe` to open the Registry Editor.
    - b. In the Registry Editor, locate the key, `HKEY_LOCAL_MACHINE\SOFTWARE\Candle`.
    - c. Right-click the key and click **Permissions**.
    - d. 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.
    - e. Click the user ID in the list.
    - f. 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**.
    - g. In the Registry Editor, locate the key, `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib`.
    - h. Right-click the key and click **Permissions**.



- i. 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.
- j. Click the user ID in the Group or user names list.
- k. 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**.
- l. Close the Registry Editor.
- m. Locate the *install\_dir* directory, where *install\_dir* is the agent installation directory.
- n. Right-click the directory and click **Properties**.
- o. 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.
- p. Click the user ID in the Group or user names list.
- q. 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.
- r. 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 423.

## Configuring IBM MQ (WebSphere MQ) for data enablement

Before you configure the IBM MQ(formerly WebSphere MQ) agent, it is recommended to the configure 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 423.
- 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 424.

### 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 422.
- 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 423.
- 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 **dsqmq** 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 425.
- To configure the agent without interaction, edit the silent response file and then run the configuration script. See [“Silent configuration”](#) on page 425.

## 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*.

4. 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.

5. 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:

1. 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.

2. Required: For **QMNAME**, specify the name of the queue manager to be monitored.
3. 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*.

4. 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.
5. 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"
```

6. 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 124](#).

---

## Chapter 11. 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.

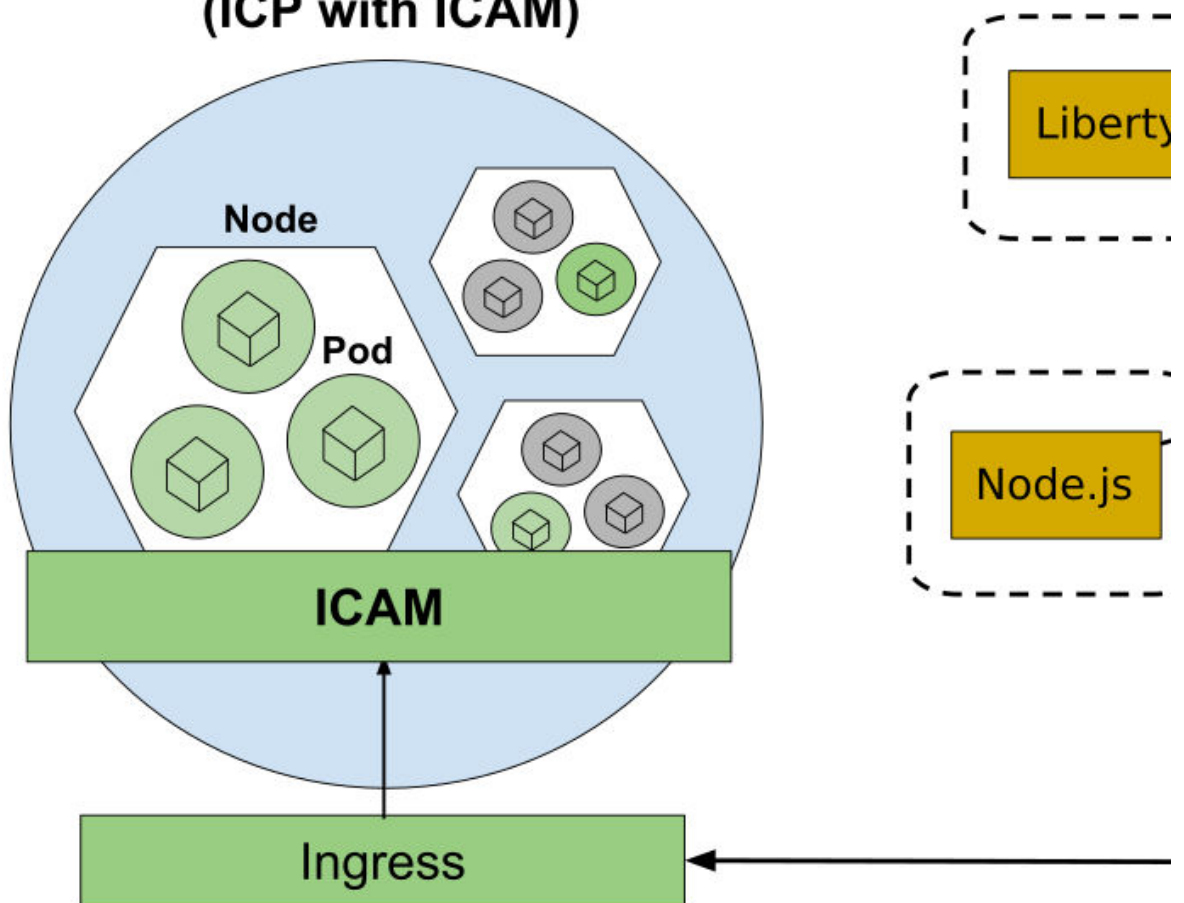
- The Kubernetes data collector is not embedded into the application workloads: it is installed outside the container workloads and typically collects the performance metrics by connecting to the workload's management interface. Kubernetes workloads can be monitored by the Kubernetes data collector.
- 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)

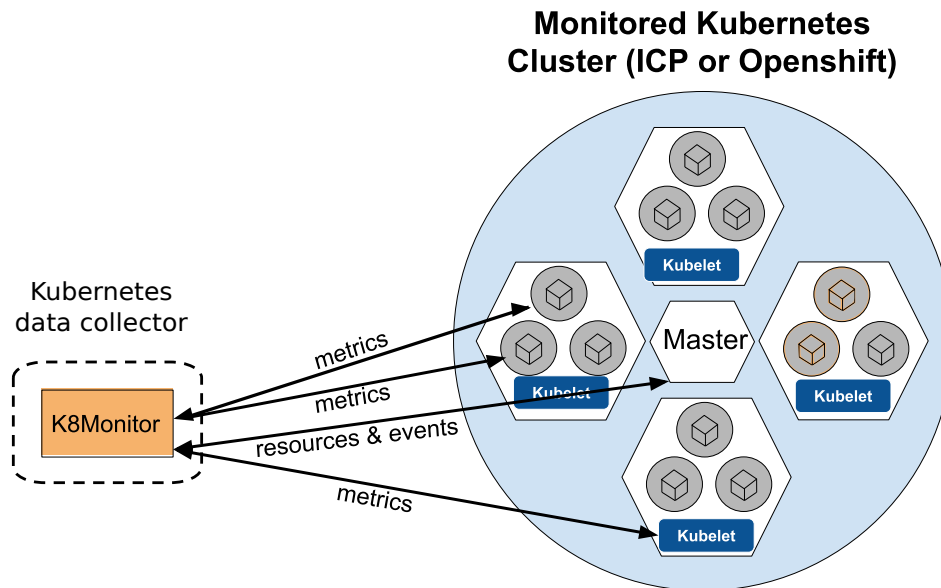


## Kubernetes data collector

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 434
- Helm client and server (Tiller) version 2.11.0 or higher. Otherwise, follow the instructions in [“Configuring Kubernetes monitoring without Helm”](#) on page 437
- 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 58, your monitored environment can be either Size0 or Size1. For Kubernetes monitoring, Size0 is for monitoring up to 250 containers, Size1 for up to 1000 containers, and Size1+ for larger environments:

*Table 56. Sizings for Kubernetes monitoring*

Description	Size0	Size1	Size1+
Monitored containers	250	1000	each additional 1000 containers
Resource Event Collector CPU (millicores)	250	1000	1000
Resource Event Collector Memory (Mi)	350	350	50
Metric Collector CPU (millicores)	500	2000	2000
Metric Collector Memory (Mi)	350	450	100
Total CPU (millicores)	750	3000	3000
Total Memory (Mi)	700	800	150

- For each additional 1000 containers beyond the first 1000 (Size1+), you must extract `app_mgmt_k8sdc_helm.tar.gz`, specify the additional resources, and compress `app_mgmt_k8sdc_helm.tar.gz` again before you run the Ansible script in step “9” on page 432.

#### 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?](#)

#### 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 CC3FMEN) from [IBM Passport Advantage](#).  
For more information, see [“Part numbers” on page 53](#).
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 431](#) and move `ibm-cloud-apm-dc-configpack.tar` to your working directory:

```
tar -xvf appMgtDataCollectors_2019.3.0.tar.gz
cd app_mgmt_k8sdc
tar -xvf app_mgmt_k8sdc.tar.gz
cd app_mgmt_k8sdc
mv my_path_to_configpack/ibm-cloud-apm-dc-configpack.tar
```

8. Extract `app_mgmt_k8sdc.tar.gz` and edit `values.yaml` to specify `environmentSize`:

```
tar xvf app_mgmt_k8sdc.tar.gz --warning=no-timestamp
sed -i 's/environmentSize:.*environmentSize: "size1"/' k8monitor/values.yaml
mv app_mgmt_k8sdc.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 [Table 56 on page 430](#)).

Before running the installation script in the next step, you must decide on important deployment configurations ([Table 57 on page 432](#)) to be passed to the install script.

9. Use [Table 57 on page 432](#) 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 57. Ansible-playbook configuration options*

Configuration option	Description	Required	Default configuration value
<code>cluster_name</code>	Unique name to distinguish your cluster from the other clusters being monitored. Use alphanumeric characters only, with no spaces. If you enter invalid characters, they are removed from the name. The assigned cluster name in the example is <code>myCluster</code> . If you want to change the <code>cluster_name</code> after deployment, see "What to do next" later in this topic.	No	UnnamedCluster
<code>release_name</code>	The Helm release name for the Kubernetes data collector. Choose a release name that does not yet exist in your environment.	No	<code>icam-kubernetes-resources</code>
<code>namespace</code>	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	<code>default</code>
<code>docker_registry</code>	The host and port of the Docker registry where you want to store the data collector images. Example: <code>mycluster.icp:8500</code> .	No	<code>mycluster.icp:8500</code>

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:

10. 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 “Validating the Kubernetes installation logs” on page 440). The data collector startup creates a Kubernetes event, which generates an informational incident.

11. 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 611. 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 589.

## 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.

## What to do next

- For each Kubernetes cluster that you want to monitor, repeat the steps starting at step “3” on page 431 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 104.

- 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-k8sdc-k8monitor`.
- For troubleshooting the deployment, see [“Kubernetes data collector issues”](#) on page 446.

**Important:** The `ibm-k8monitor-config` ConfigMap is created in your default namespace as part of Kubernetes data collector deployment. Do not delete, move, or rename this resource. 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. 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.

## 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 437
- 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=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?](#).

### 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 CC3FMEN) from [IBM Passport Advantage](#).

For more information, see [“Part numbers”](#) on page 53.

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_2019.3.0.tar.gz
cd app_mgmt_k8sdc
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_2019.3.0` 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
```

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
```

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 \  

```


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-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 [“Validating the Kubernetes installation logs”](#) on page 440). 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 611. 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 589.

## 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.

## 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 104](#).
- 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-k8sdc-k8monitor`.
- For troubleshooting the deployment, see [“Kubernetes data collector issues” on page 446](#).

**Important:** The `ibm-k8monitor-config` ConfigMap is created in your default namespace as part of Kubernetes data collector deployment. Do not delete, move, or rename this resource. 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. 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.

## 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.

### 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 CC3FMEN) from [IBM Passport Advantage](#).



For more information, see “Part numbers” on page 53.

2. Extract the Docker images:

```
tar xvf appMgtDataCollectors_2019.3.0.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_imagetag
docker push my_clustername:my_clusterport/my_namespace/k8-monitor:my_imagetag
```

b) Load the k8sdc-operator Docker image to the repository:

```
docker load -i app_mgmt_k8sdc/app_mgmt_k8sdc_operator_docker.tar.gz
docker tag my_repotags my_cluster_ca_domain:my_docker_registry_port/my_namespace/k8sdc-
operator:my_imagetag
docker push my_cluster_ca_domain:8500/my_namespace/k8sdc-operator:my_imagetag
```

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\_imagetag* 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/ibmcloudappmgmt_v1alpha1_k8sdc_crd.yaml
```

- b) Apply any necessary value changes to the "Custom Resource" for operator spec by editing the `deploy/crds/ibmcloudappmgmt_v1alpha1_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/ibmcloudappmgmt_v1alpha1_k8sdc_cr.yaml
```

- c) Create the service-account for the operator:  
where

```
my_cluster_ca_domain is the target cluster CA domain to monitor
my_namespace is the target namespace on the cluster
my_imagetag is the Docker image tag, which you can get from the RepoTags in the
manifest.json within the archive file
```

```
kubectl create -f deploy/service_account.yaml
```

- d) Apply the **imagePullSecrets** that you created in step "5" on page 438 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 “Validating the Kubernetes installation logs” on page 440). 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 611. 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 589.

## 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.

## 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 104.
- 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-k8sdc-k8monitor`.
- For troubleshooting the deployment, see “Kubernetes data collector issues” on page 446.

**Important:** The `ibm-k8monitor-config` ConfigMap is created in your default namespace as part of Kubernetes data collector deployment. Do not delete, move, or rename this resource. 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. 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.

## Validating the Kubernetes installation logs

As a best practice, review the installation logs to ensure that the Kubernetes data collector has been properly configured.

### Procedure

Stream the K8Monitor logs and ensure that all initial checks have passed:

```
kubectl logs -f my_ReleaseName-k8monitor --namespace=my_Namespace
```

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 may have provided invalid configuration parameters. Before retrying 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 -namespace=my_namespace delete deployment my_releasename-k8monitor
```

- b) Delete the secrets:

```
kubectl -namespace=my_namespace delete secret dc-secret
kubectl -namespace=my_namespace delete secret ibm-agent-https-secret
kubectl -namespace=my_namespace delete secret my_releasename-adminintnants
```

- 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 Helm, enter these kubectl commands:

```
kubectl delete -f deploy/crds/ibmcloudappgmt_v1alpha1_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/ibmcloudappmgmt_v1alpha1_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.ibmcloudappmgmt.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.ibmcloudappmgmt.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.ibmcloudappmgmt.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 445](#) 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 [“Validating the Kubernetes installation logs” on page 440](#)), 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 447](#) 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 162](#).)

## Runtime data collectors

Runtime data collectors monitor Node.js, Liberty, J2SE, and Python applications and are embedded within the application workloads.

### Common topics

Some topics are common to Node.js, Liberty, J2SE, and Python 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
```

**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 data collector. By default, the *service\_account\_name* is default.

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
```

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:  
  - ""  
  resources:  
  - nodes
```

```
- services
- configmaps
- pods
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 457](#), [“Monitoring Node.js applications in Kubernetes environment” on page 451](#), [“Monitoring J2SE applications in Kubernetes environment” on page 465](#), and [“Monitoring Python applications in Kubernetes environment” on page 472](#).

## 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.

System requirements:

- IBM Cloud Private 3.2.0 and above, or OpenShift 3.11
- The supported operating systems and platforms are the same as IBM Cloud Private and OpenShift support. For more information about IBM Cloud Private system requirements, see [IBM Cloud Private Knowledge Center](#). For more information about OpenShift system requirements, see [OpenShift documentation](#).
- Node.js
  - 8.0.0 and future fix packs
  - 10.0 and future fix packs
  - 12.0 and future fix packs

The Node.js data collector helps you to manage the performance and availability of the following:

- Node.js and Microclimate-based applications in IBM Cloud Private

### 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:
  - a. To update the `package.json`, add `"appmetrics": "^4.0.0"` as a dependency.
  - b. 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:

- a. Unpack the data collectors package according to your Node.js Runtime version, for example, for `appMgtDataCollectors_2019.3.0.tar.gz`:

```
tar xzf appMgtDataCollectors_2019.3.0.tar.gz
cd appMgtDataCollectors_2019.3.0
tar xzf app_mgmt_runtime_dc_2019.3.0.tar.gz
cd app_mgmt_runtime_dc_2019.3.0
tar xzf nodejs_datacollector_2019.3.0.tgz
tar xzf ibmapm-greenfield-v8-1x64.tgz
```

For more information, see [“Obtaining the server configuration information” on page 449](#).

- b. 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
```

- c. 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 448](#).

Check that you downloaded the configuration package to obtain the server information, for more information, see [“Obtaining the server configuration information” on page 449](#).

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 449](#), 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 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 448.](#)

Check that you downloaded the configuration package to obtain the server information, for more information, see [“Obtaining the server configuration information” on page 449.](#)

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 451.](#) 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 449.](#)

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_2019.3.0.tgz`, and `ibm-cloud-apm-dc-configpack.tar` file that you download from the server. `apply_configpack.sh`

### 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 449,](#) 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 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 Node.js data collector enables OpenTracing function. You can disable OpenTracing by setting the environment variable to false.  Example: <pre>- name: OPENTRACING_ENABLED   value: false</pre> <b>Note:</b> If you install Node.js data collector by requiring "appmetrics", the OpenTracing is disabled by default.
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.	When the OpenTracing function is enabled, you can set the OpenTracing sampler rate. Example: <pre>- name: OPENTRACING_SAMPLER   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 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 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 WebSphere Liberty profile.

System requirements:

- IBM Cloud Private 3.2.0 and above, or OpenShift 3.11
- The supported operating systems and platforms are the same as IBM Cloud Private and OpenShift support. For more information about IBM Cloud Private system requirements, see [IBM Cloud Private Knowledge Center](#). For more information about OpenShift system requirements, see OpenShift documentation.
- WebSphere Application Server Liberty Core
  - 8.5.5 and future fix packs
  - 16.0.0.2
  - 16.0.0.3
  - 16.0.0.4
  - 17.0.0.1
  - 17.0.0.2
  - 18.0.0.2
  - 19.0.0.8

The Liberty data collector helps you to manage the performance and availability of the following:

- Java-based microservices or Liberty applications in IBM Cloud Private.

You can configure the data collector to send data to the Cloud App Management server.

### (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_2019.3.0.tar.gz
cd appMgtDataCollectors_2019.3.0
tar xzf app_mgmt_runtime_dc_2019.3.0.tar.gz
cd app_mgmt_runtime_dc_2019.3.0
```

### 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 448.

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 456.

Check that you downloaded the configuration package to obtain the server information, for more information, see [“Obtaining the server configuration information”](#) on page 449.

#### 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 449, 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 Liberty application. You must have the write access to the server folder.

- If your environment has access to the WebSphere Liberty repository where the Liberty application container is built, then add the following to your Docker file:

```
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 your environment doesn't have access to the WebSphere Liberty repository where the Liberty application container is built, then add the following to your Docker file:

```
RUN chmod 777 liberty_server_dir
COPY path_to_es_a_file /opt/
RUN 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_es_a_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 path to the 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. If you want to update this file, see the following Liberty Data Silent Config file:

```
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.

**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. 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
```

4. Open your application deployment yam1 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” on page 457.

If you are working with a local application deployment yam1, you must run the following command for your changes to take effect:

```
kubectl create -f application_deployment_yam1_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 Docker file 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 Docker file:

```
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/lx8266 \

  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. To enable your application with the Cloud App Management data collector, then just add the following to the Docker file:

```
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 457 and step "4" on page 459 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 449](#).
- Check whether you downloaded the Liberty data collector package. For more information, see ["\(Conditional\) Downloading the Liberty data collector" on page 456](#).

## 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.

- 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 Docker file and add the following:

```
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) Add the following lines to the Docker file of your Liberty application:

```
COPY path_to_esa_file /opt/  
RUN installUtility install --acceptLicense /opt/  
javametrics.liberty.icam-1.2.1.esa  
COPY path_to_silent_file /opt/ibm/wlp/usr/extension/liberty_dc/bin/  
COPY path_to_configpack_file /opt/  
RUN /opt/ibm/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.
- *path\_to\_configpack\_file* is the relative path for the configuration package file to the current directory. For example, /tmp/ibm-cloud-apm-dc-configpack.tar.

d) Build the new docker image.

```
docker build -t <application image name>:<image tag>
```

e) 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.

Table 59. Supported user-defined environment variables for Liberty monitoring

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>

### Configuring digital experience monitoring (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 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 429 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 429.
- 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 457.
- 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,

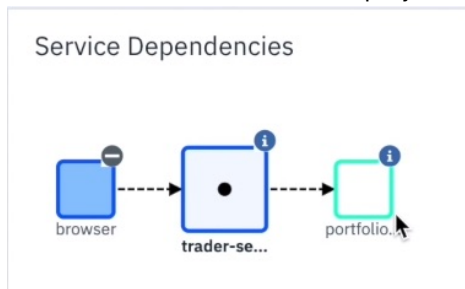
```
- name: JAEGER_SAMPLER_TYPE
  value: "ratelimiting"
- name: JAEGER_SAMPLER_PARAM
  value: 10
```

For more information, see “[Customizing the Liberty data collector](#)” on page 462.

### 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.

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.

### 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:
  - 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) Roll back changes to the Docker file by removing the following lines:

```
COPY /javametrics.liberty.icam-1.2.1.esa /opt/
installUtility install --acceptLicense /opt/javametrics.liberty.icam-1.2.1.esa
COPY /silent_config_liberty_dc.txt /opt/ibm/wlp/usr/extension/liberty_dc/bin/
RUN /opt/ibm/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent -javametrics
```

- c) 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:

- a) Roll back changes to the Docker file by removing the following lines:

```
COPY /javametrics.liberty.icam-1.2.1.esa /opt/
installUtility install --acceptLicense /opt/javametrics.liberty.icam-1.2.1.esa
COPY /silent_config_liberty_dc.txt /opt/ibm/wlp/usr/extension/liberty_dc/bin/
RUN /opt/ibm/wlp/usr/extension/liberty_dc/bin/config_unified_dc.sh -silent -javametrics
```

- b) Rebuild the Docker images and redeploy the application.

- If you deploy your Liberty data collector on bare-metal node or VMs, disable Liberty data collector, stop Liberty server, and then run the following command:

```
./unconfig_liberty_dc.sh [-silent <silent_file>]
```

## 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.

System requirements:

- IBM Cloud Private 3.2.0 and above, or OpenShift 3.11
- The supported operating systems and platforms are the same as IBM Cloud Private and OpenShift support. For more information about IBM Cloud Private system requirements, see [IBM Cloud Private Knowledge Center](#). For more information about OpenShift system requirements, see [OpenShift documentation](#).
- IBM Runtime Environment, Java Technology Edition 8, and future fix packs
- Oracle Java SDK/JRE/JDK 8.0 and future fix packs

The J2SE data collector helps you to manage the performance and availability of Spring Boot based applications and other stand-alone Java applications in IBM Cloud Private.

You can configure the J2SE data collector to send data to the Cloud App Management server.

### 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.tar.gz`) by running the following command:

```
tar -xzf appMgtDataCollectors_2019.3.0.tar.gz
cd appMgtDataCollectors_2019.3.0
tar -xzf app_mgmt_runtime_dc_2019.3.0.tar.gz
cd app_mgmt_runtime_dc_2019.3.0
```

### 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 448](#).
- Check whether the J2SE data collector is available for download from Passport Advantage. For more information, see [“Downloading the J2SE data collector” on page 465](#).
- Check whether you downloaded the configuration package to obtain the server information. For more information, see [“Obtaining the server configuration information” on page 449](#).

### 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 449, 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 -xvf 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=FALSE
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 yml 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 yml 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 449](#).
- Check whether you downloaded the J2SE data collector package. For more information, see [“Downloading the J2SE data collector” on page 465](#).

#### Procedure

1. Extract the j2se\_datacollector.tgz file that you get from [“Downloading the J2SE data collector” on page 465](#).

```
tar -xf 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 449](#), 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.

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 -xf /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.

Table 60. Supported user-defined environment variables for J2SE monitoring

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:  - name: OPENTRACING_ENABLED value: false
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:  - name: JAEGER_SAMPLER_TYPE value: <i>probabilistic</i> - name: JAEGER_SAMPLER_PARAM value: "0.1"
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:  - name: LATENCY_SAMPLER_PARAM value: "0.2"

### 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 yml 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 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 IBM Cloud Private
- Local Python applications with Django or Flask frameworks

**Important:** The Python data collector can monitor only internal web servers of Django or Flask frameworks.

### System requirements

- IBM Cloud Private 3.2.0 and above, or OpenShift 3.11
- The supported operating systems and platforms are the same as IBM Cloud Private and OpenShift support. For more information about IBM Cloud Private system requirements, see [IBM Cloud Private Knowledge Center](#). For more information about OpenShift system requirements, see [OpenShift documentation](#).
- Python 2.7, 3.6.x, and 3.7.x
  - Note:** For Python 2.7 series, it is required to install Python 2.7.15 and above.
- Django 1.4.10 or above
- Flask 1.0.x and 1.1.x or above

### 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_2019.3.0.tar.gz
cd appMgtDataCollectors_2019.3.0
tar xzf app_mgmt_runtime_dc_2019.3.0.tar.gz
cd app_mgmt_runtime_dc_2019.3.0
```

## 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 448.
- Ensure that you downloaded the configuration package to obtain the server information. For more information, see [“Obtaining the server configuration information”](#) on page 449.
- 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 471.

### 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. 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,
})

.....
.....
```

3. Go to the `ibm-cloud-apm-dc-configpack` directory where you extract the configuration package in “Obtaining the server configuration information” on page 449, 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.

4. Update the application yaml 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
```

5. Build the new Docker image.
6. Update the application yml 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 474.
- 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 475.

#### *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 449.
- Check whether you downloaded the Python data collector package from Passport Advantage. For more information, see [“Downloading the Python data collector”](#) on page 471.

#### 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. 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. 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 449.
- 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 471.

#### **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. 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,
})

.....
.....
```

## 3. Build the new Docker image.

## 4. 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 yml file.

For on-premises applications, set the following variables as environment variables or in the `config.properties` file.

Table 61. 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 61. 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 61. 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 475.

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 yaml 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.

## 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.

There are four different types of synthetic tests you can create.

### REST API

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 483.

### Scripting REST API

Test and monitor a number of REST APIs in a sequence. For more information, see [“Creating a Scripting REST API synthetic test”](#) on page 486.

### Webpage

Test a single web page for availability and browser response time. For more information, see [“Creating a web page synthetic test”](#) on page 490.

### Selenium script

Test simulated user behavior. For more information, see [“Creating a Selenium script test”](#) on page 491.

## Installing a 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_2019.3.0.tar.gz` from Passport Advantage. This contains a sub package that is called `app_mgmt_syntheticpop_xlinux.tar.gz`, this 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 640

### Procedure

1. Download and unpack the data collectors installation eImage from Passport Advantage (`appMgtDataCollectors_2019.3.0.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 132.
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_2019.3.0.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

4. To make sure that the prerequisites are met to install the Synthetics PoP in your environment, run the following pre-check script.

```
./pre-check.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-clouddc-configpack.tar
```

You are prompted to enter the following parameters:

Table 62. 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 62. 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-clouddc-configpack.tar` as described in step “5” on page 481. 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 Private Locations section in the Synthetic test editor page, click the delete button.

#### Collect logs

Run the following command to collect the log files from the Synthetics PoP container: `./pdcollect.sh`

## Results

You can now log in to Cloud App Management console, and create a synthetic test. The Synthetics PoP is shown as a Private Location based on the name value you specified.

## 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.

Request

2. 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.
3. 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 name** 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 user name 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.

4. 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.

Select one of the following:

#### 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.

For each condition, enter a property to test for in the **Target** field, and a value to test for in the **Value** field. Select an operator from the **Operation** drop-down menu. Finally, choose an **Event severity** of **Warning** or **Critical** for your condition. For examples, see [“Customizing response validation” on page 484](#).

Response validation

5. Configure the warning and critical events conditions for your synthetic test in the **Response Validation** section. Accept the defaults, or edit the **Value** and **Unit** for each row. Response times that exceed 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 the Alert Triggers step.

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 496](#).

Verify

6. 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

### 7. 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.

### 8. The **Private 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 Private Location, see [“Installing a Synthetics PoP” on page 480](#).

## Alert 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 **Failure detected** to **Off**. To increase the number of failures allowed before a critical alert is triggered, select **consecutively** under **Failure detected**. The default number of consecutive failures is 2, but this number can be customized.

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, select **consecutively** under **Threshold breached**. The default number of slow response times is 2. This number can be customized.

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 496](#).

## 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 Synthetics 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. Select a node . js test script. For information about creating a node . js script to test, see [“Create a REST API test case”](#) on page 487.

Response validation

4. Configure the warning and critical events conditions for your synthetic test in the **Response Validation** section. Accept the defaults, or edit the **Value** and **Unit** for each row. Response times that exceed 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 the Alert Triggers step.

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 496.

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 **Private 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 Private Location, see [“Installing a Synthetics PoP”](#) on page 480.

## 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 **Failure detected** to **Off**. To increase the number of failures allowed before a critical alert is triggered, select **consecutively** under **Failure detected**. The default number of consecutive failures is 2, but this number can be customized.

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, select **consecutively** under **Threshold breached**. The default number of slow response times is 2. This number can be customized.

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 496.

## 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");
  }
);
});

```

- 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()`.

- 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");
```

- Use the following syntax for a GET request:

```

let baseUrl = 'http://localhost:3000/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();
  });
});

```

- 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 = 'http://localhost:3000';
describe('post json content',function(complete){
  // Define endpoint URL
  let url = baseUrl + "/messages";
  // Define JSON data
  let data = {"text": "Hi LiLi is here"};
  // Define headers
  let header1 = {"context-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 == 200, "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. Accept the defaults, or edit the **Value** and **Unit** for each row. Response times that exceed 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 the Alert Triggers step.

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 496](#).

##### Blocking and filtering

5. In the Blacklist field, enter any URL 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.

##### 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 **Private 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 Private Location, see [“Installing a Synthetics PoP” on page 480](#).

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 **Failure detected to Off**. To increase the number of failures allowed before a critical alert is triggered, select **consecutively** under **Failure detected**. The default number of consecutive failures is 2, but this number can be customized.

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, select **consecutively** under **Threshold breached**. The default number of slow response times is 2. This number can be customized.

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 496](#).

## 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 492](#).

#### Response validation

4. Configure the warning and critical events conditions for your synthetic test in the **Response Validation** section. Accept the defaults, or edit the **Value** and **Unit** for each row. Response times that exceed 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 the Alert Triggers step.

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 496](#).

#### Blocking and filtering

5. In the Blacklist field, enter any URL 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.

#### Authentication

6. If the web page you are testing requires authentication (NTLM or basic), enter the username and password.

**Note:** The Verify button is not available for Selenium synthetic tests. The Verify button will bypass Selenium script verification if Cloud App Management server is deployed in power platform.

Review and Finish

7. 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

8. The **Private 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 Private Location, see [“Installing a Synthetics PoP” on page 480](#).

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 **Failure detected** to **Off**. To increase the number of failures allowed before a critical alert is triggered, select **consecutively** under **Failure detected**. The default number of consecutive failures is 2, but this number can be customized.

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, select **consecutively** under **Threshold breached**. The default number of slow response times is 2. This number can be customized.

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 496](#).

### **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 495](#).

### 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 64. 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 492](#).
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 492](#) 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 65. Sample script for opening the IBM Marketplace page ( <i>load_homepage.side</i> )		
Command	Target	Value
open	/	
verifyTitle	IBM Cloud Application Performance Management	

Table 66. 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 67. 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	

- 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.

## 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=("//ul[@id='select2-country-results']/li[text() = 'United States'])</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 484](#).

## Viewing Synthetic test results

Use the **Synthetic results** tab in the Cloud App Management console to visualize your synthetic API tests.

### Before you begin



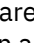




The **Synthetic results** tab is available with the IBM Cloud App Management advanced offering.

### 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 480](#).

### 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 480](#).
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.

# Chapter 12. Unified Agent

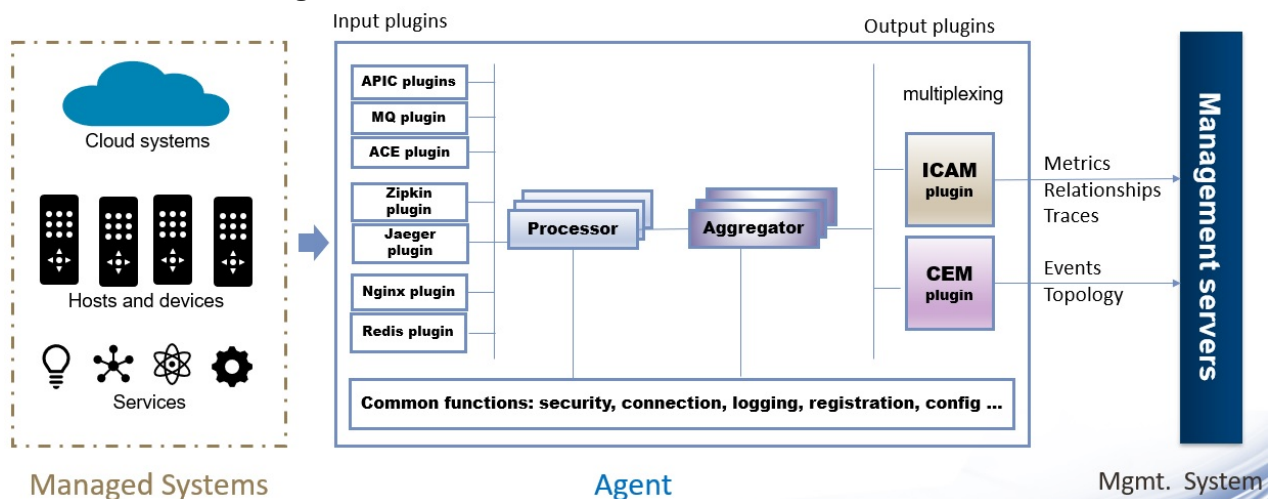
The Unified Agent is an agent for collecting, processing, aggregating, and writing metrics to your IBM Cloud App Management environment. It is based on Telegraf and creates a framework of plug-ins to provide common functions.

## 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.

### Unified Agent Architecture

Unified Agent is based on open source technology called Telegraf. The following picture shows the architecture of Unified Agent.



### Supported 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

## 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

- Kubectl client on the environment from where you are installing
- Kubernetes version 1.7 or higher (available with IBM Cloud Private version 3.1.0)
- IBM Cloud Private 3.1.0 or higher environment that you want to monitor

### Special prerequisites for plug-ins

<i>Table 68. Prerequisites for each plug-in of Unified Agent</i>			
<b>Plug-in</b>	<b>Supported environment</b>	<b>Supported version for monitored target</b>	<b>Other requirements</b>
Jaeger and Zipkin	—	—	—
NGINX	—	NGINX version 1.12 or higher	Make sure the NGINX monitoring interface is enabled. For more information, see <a href="#">“Enabling the NGINX monitoring interface” on page 502.</a>
Redis	—	Redis version 3.2 or higher	—
IBM API Connect	<ul style="list-style-type: none"> <li>• Kubernetes</li> <li>• IBM Cloud Private</li> <li>• OpenShift</li> </ul> <p><b>Note:</b> For specific version information, see <a href="#">IBM API Connect Knowledge Center.</a></p>	IBM API Connect 2018.x	<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” on page 505.</a></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>

<i>Table 68. Prerequisites for each plug-in of Unified Agent (continued)</i>			
<b>Plug-in</b>	<b>Supported environment</b>	<b>Supported version for monitored target</b>	<b>Other requirements</b>
IBM App Connect Enterprise	IBM Cloud Private 3.2.1	<ul style="list-style-type: none"> <li>IBM App Connect Enterprise 11.0.0.5-amd64</li> <li>IBM App Connect Enterprise 11.0.0.5</li> </ul>	<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 505.</li> <li>Make sure the Prometheus service for IBM App Connect Enterprise is launched.</li> </ul>
IBM MQ	IBM Cloud Private 3.2.1	IBM MQ Advanced certified container 4.1.1  For more information, see <a href="#">IBM MQ Knowledge Center</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 505.</li> <li>Make sure the Prometheus service for IBM MQ is launched.</li> </ul>

### **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_2019.3.0.tar.gz` from [IBM Passport Advantage](#).

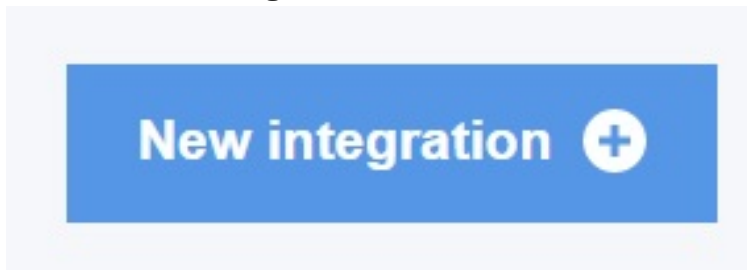
For more information, see “Part numbers” on page 53.

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.

A file with name `ibm-cloud-apm-dc-configpack.tar` 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>
> kubectrl 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      10.1.13.68      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
```

## Deploying 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 499.

### About this task

#### Procedure

1. Extract the `unifiedAgent_2019.3.0.tar.gz` package to get cloud monitoring media.

```
tar xzf unifiedAgent_2019.3.0.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, and IBM MQ by the option `-p`.

```
USAGE: deploy.sh
[ -p Plugin name ] # e.g -p opentracing:nginx:redis:ibmmq:ibmapic:ibmace
[ -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 -p opentracing:ibmmq:ibmapic:ibmace:nginx:redis -r monitor
-c /tmp/ibm-cloud-apm-dc-configpack.tar
```

### Important:

- 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 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 508](#).
  - To configure the NGINX plug-in, see [“Configuring NGINX monitoring in Unified Agent ” on page 509](#).
  - To configure the Redis plug-in, see [“Configuring Redis monitoring in Unified Agent ” on page 509](#).
  - To configure the IBM APIC plug-in, see [“Configuring IBM API Connect monitoring in Unified Agent ” on page 510](#).
  - To configure the IBM ACE plug-in, see [“Configuring IBM App Connect Enterprise monitoring in Unified Agent ” on page 510](#).
  - To configure the IBM MQ plug-in, see [“Configuring IBM MQ monitoring in Unified Agent ” on page 511](#).
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, 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:

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 [“Deploying the Unified Agent”](#) on page 506 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 [“Deploying the Unified Agent”](#) on page 506 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 [“Deploying the Unified Agent”](#) on page 506 to deploy the Unified Agent.

2. Specify the Redis service address that you want to monitor, for example, `tcp://:redisPassw0rd@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 [“Deploying the Unified Agent” on page 506](#) 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 [“Deploying the Unified Agent” on page 506](#) 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 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 [“Deploying the Unified Agent” on page 506](#) 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.

## 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.
- `ualk-ibmapic` 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 <configmap> -n <namespace>
```

To reconfigure each plug-in in Unified Agent, open the configmap `cloud-ya-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.

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
```

**Results**

The Unified Agent is removed from your system.



---

## Chapter 13. 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.

The following V6 agents are supported to connect to the Cloud App Management server.

- Cisco UCS agent
- DataPower agent
- Db2 agent
- HTTP Server agent
- IBM Integration Bus agent
- JBoss agent
- Linux OS agent
- Linux KVM agent
- Microsoft Hyper-V Server agent
- Microsoft IIS agent
- Microsoft SQL Server agent
- NetApp Storage agent
- Oracle Database agent
- UNIX OS agent
- VMware VI agent
- WebSphere Applications agent
- IBM MQ(formerly WebSphere MQ) agent
- Windows OS agent

#### 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:


- Both HTTP and HTTPS connections to the Cloud App Management server are supported.

- 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 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.
- **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 518.
- 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 521.

#### **Locally configuring the agent to connect to Cloud App Management server**

##### Procedure

1. Extract the agent patch to the local system where the V6 agent is installed.
  - For HTTP connection, extract 6.3.0.7-TIV-ITM\_TEMA-IF0003.tar or 6.3.0.7-TIV-ITM\_TEMA-IF0003.
  - For HTTPS connection, extract 6.3.0.7-TIV-ITM\_TEMA-IF0008.tar or 6.3.0.7-TIV-ITM\_TEMA-IF0008.

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, 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](#).

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 524.
- (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 to a temporary directory.
  - For HTTP connection, extract `6.3.0.7-TIV-ITM_TEMA-IF0003.tar` or `6.3.0.7-TIV-ITM_TEMA-IF0003`.
  - For HTTPS connection, extract `6.3.0.7-TIV-ITM_TEMA-IF0008.tar` or `6.3.0.7-TIV-ITM_TEMA-IF0008`.

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, 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.

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 521.

- 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](#).

- Download the agent configuration pack from the Cloud App Management console. The downloaded package contains agent configuration files for server connection.
  - Log in to the Cloud App Management console and click **Get Started**.
  - Click **Administration > Integrations > Configure an integration (Incoming)**.
  - In the Standard monitoring agents section, go to the **ITM/ITCAM Agents** tile and click **Configure**.
  - Click **Download file** to download the `ibm-cloud-apm-v6-configpack.tar` file.
- 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.

- 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 524.
- (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 528 \(bn\)](#)

[“Db2 agent” on page 528 \(ud\)](#)

[“Hyper-V Server agent” on page 528 \(hv\)](#)

[“IBM Integration Bus agent” on page 529\(qi\)](#)

[“Linux OS agent” on page 529 \(lz\)](#)

[“Microsoft IIS agent” on page 529 \(q7\)](#)

[“Microsoft SQL Server agent” on page 529 \(oq\)](#)

[“UNIX OS agent” on page 530 \(ux\)](#)

[“WebSphere Applications agent” on page 530 \(yn\)](#)

[“IBM MQ\(formerly WebSphere MQ\) agent” on page 530 \(mq\)](#)

[“Windows OS agent” on page 530 \(nt\)](#)

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.pipe:#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 524](#)

#### 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.

- Cisco UCS 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
- SAP agent
- SAP HANA Database agent
- SAP NetWeaver Java Stack agent
- Skype for Business Server agent
- Tomcat agent
- VMware VI agent
- UNIX OS 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:**

- Both HTTP and HTTPS connections to the Cloud App Management server are supported. However it is not supported to reconnect the agents to the Cloud App Management server if the connection to the Cloud App Management server uses the https protocol.
- 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 533.

## 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 -i
```

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`.

**Restriction:** If the connection to the Cloud App Management server is using the HTTPS protocol, then it is not possible to switch back to connect the V8 agents to the Cloud App Management server again.

## 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:
    - a) Click **Actions** in the menu bar on the left, and click **Create Action** to add a new template.
    - b) Select the **Make an HTTP Request** radio button, and click **OK**.
    - c) Enter a name for the action and select the template you created from the **HTTP Request Template** list.
    - d) 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": {
```



- 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 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/SSHTQ/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 548 as source.
  - “[Configure New Relic Alerts](#)” on page 548 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 547.
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 547.
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 are want to receive event information from Prometheus included in IBM Cloud Private, the following steps are different. See step “7” on page 552 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://myeventsource.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 **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.

8. 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
```

9. 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.

10. 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.

11. 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.

<i>Table 70. Mandatory CEM field error messages</i>			
<b>CEM Field</b>	<b>Messages</b>	<b>Possible Causes</b>	<b>Resolutions</b>
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.



*Table 70. Mandatory CEM field error messages (continued)*

CEM Field	Messages	Possible Causes	Resolutions
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 584](#).

[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:

```
$OMNIBHOME/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

8. 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:

```
$OMNIBHOME/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 -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 560](#) and changing the `serverPort` value.

9. 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 560, then ensure you use the port number you set there.

10. Click **Save** to save the integration in Cloud App Management.
11. 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 561.

- Create a copy of the `httpTransport.properties` file you downloaded as part of step “5” on page 560, 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
```

- Edit the `apm.props` file you downloaded as part of step “5” on page 560, and add the following property:

```
TransportFile : 'path_to_file/httpsWebhookTransport.properties'
```

- Start the probe as described in “8” on page 560, but set the `-transportfile` value to the `httpsWebhookTransport.properties` file you configured: `-transportfile /path_to_file/httpsWebhookTransport.properties`.
- 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 584.

## 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 584](#).

### 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 584.

### 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 584.

### 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 584.



## 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 584.

## 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 584](#).

### 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 584.

#### **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 14. 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 606](#).

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 71. 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 584](#).

### **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 534](#).

### **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 577](#).

## 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 576 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 Affect ed
Human Resources	-HR Application Affect ed
Payroll Application	-Payroll Applications Affected

Cancel Apply

Figure 1. 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.

Summary ▾ lookup ▾ summaryupdate in Application Lookup Table

Append to field

Prepend to field

Overwrite field

Figure 2. 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 3. Resulting enriched event

Resource affected ⓘ	
Name	Description
Application	Payroll Application

Figure 4. 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 576.

### 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 577.
  - 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 574.
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 606](#).

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 570. 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 24. 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 570.

### 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 585](#), 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 585.
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 602. 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 559.

### 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 72. 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 591](#).
2. For an example of how to start managing your incidents, see [“Starting to work with incidents” on page 595](#).



## 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)

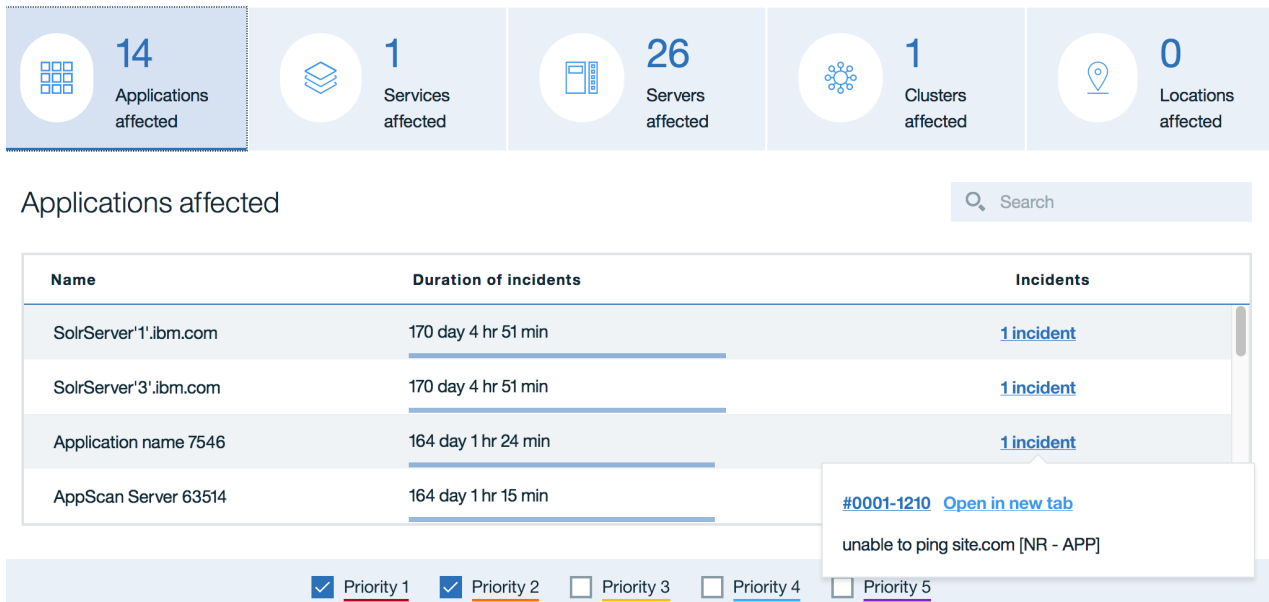


Figure 5. Resources affected by incidents widget

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)

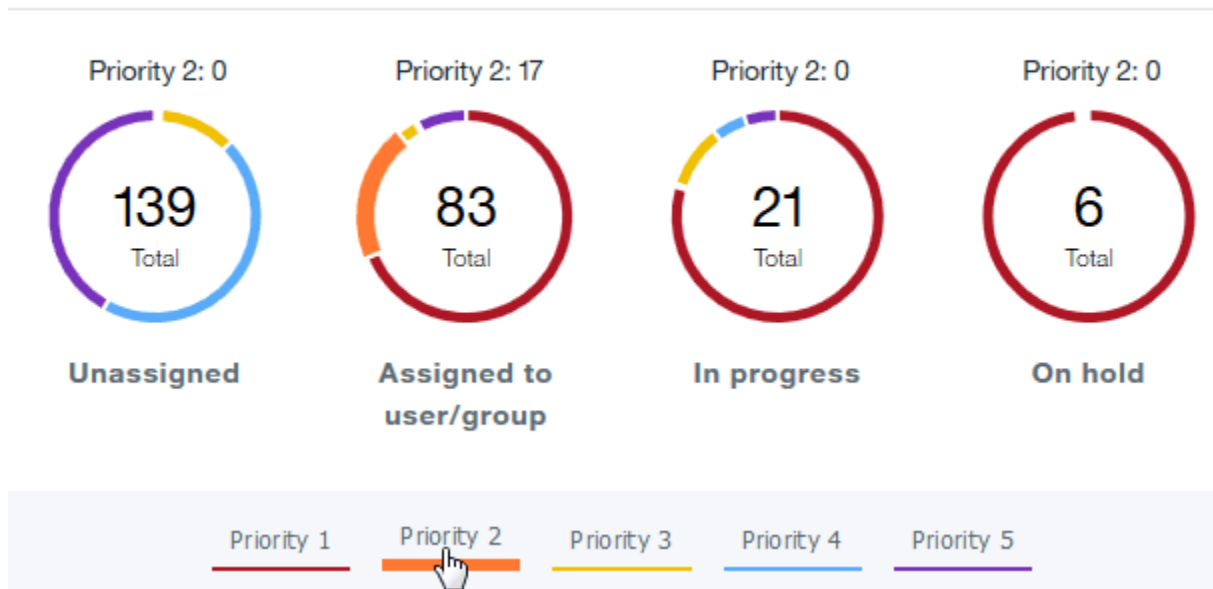


Figure 6. *State of incidents* widget

### 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.

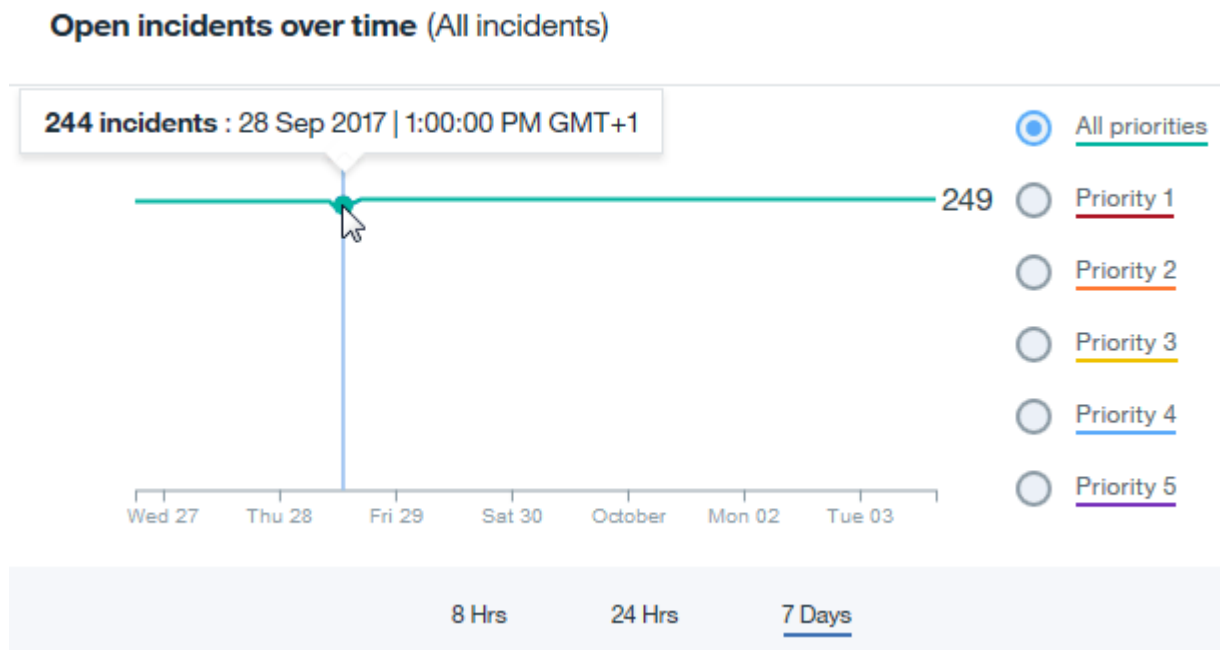


Figure 7. *Open incidents over time* widget

## 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.

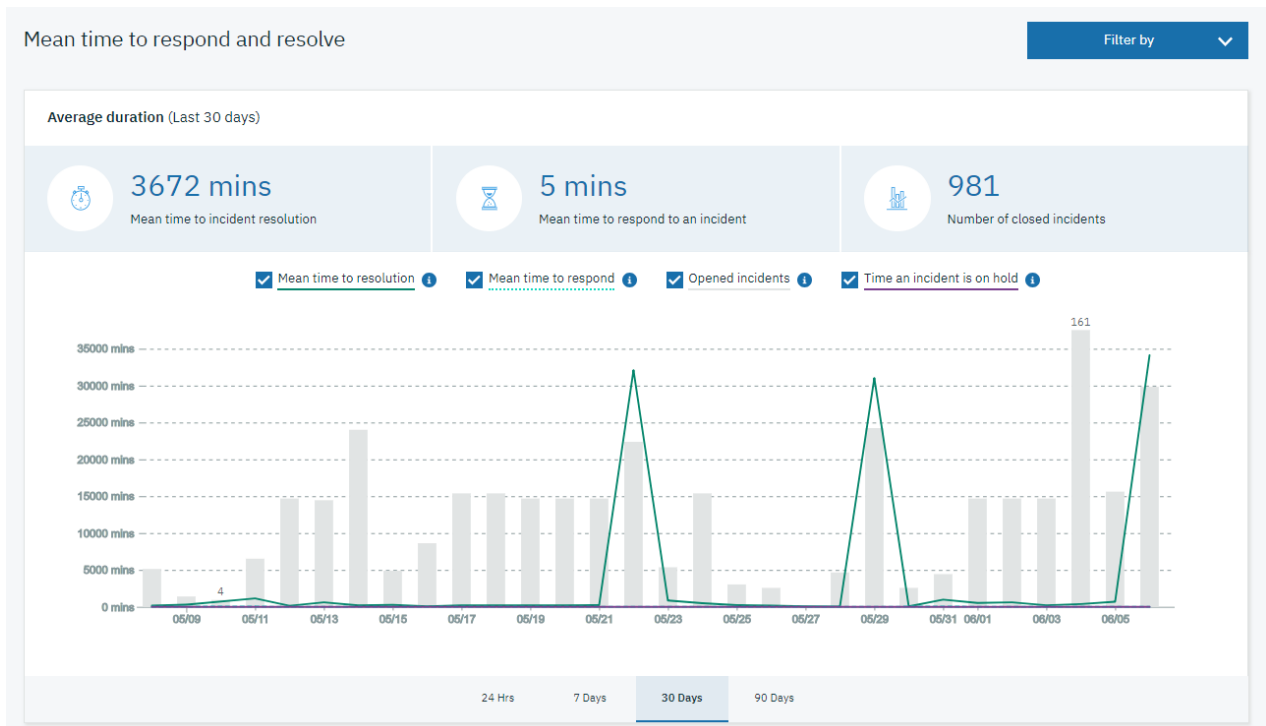


Figure 8. **Average duration** widget

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 606](#), and then set up event policies where runbooks are associated with events as described in [“Setting up event policies” on page 569](#).

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 606](#).

## 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 state when **Enable** has been turned off and the threshold is non-operational.
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 597](#).
    - To edit a threshold definition, click the threshold name. Continue to step [“4” on page 597](#).
    - 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 (or more), select its check box and click **Enable/disable** in the banner that appears. If the threshold was disabled, it is now enabled; if the threshold was enabled, it is now disabled.

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) Optional: Enter a description for the threshold.
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  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 599), 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 600](#).
8. Available only for thresholds 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.  
Examples: `echo WT_LZ_user_login is true;/scripts/clean_logs.sh &{KLZ_Disk.Disk_Name}`.
- b) The following options 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.

## Results

After you save the threshold, it starts on the resources instances that were assigned in step [“7” on page 599](#).

## 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 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
- 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 597](#).

## 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 602](#).
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 601.
<b>Email</b>	Enter the user's main email address.
<b>Role</b>	<p>Set one role for each user. The following roles can be assigned:</p> <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> <p>For more information about Cloud App Management roles, see <a href="#">“Roles”</a> on page 604.</p>
<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 604.

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 584.

## 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 604.

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

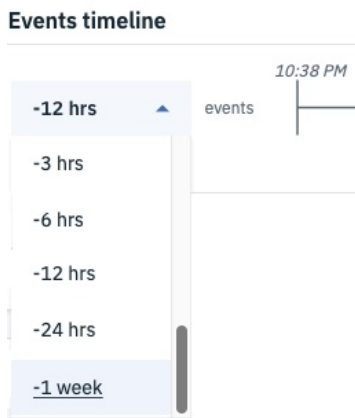
The ICAM Agents and ICAM Data Collectors are monitoring your environment for early detection of performance and availability issues.

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.





## 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 “14” on page 118 of *Installing your Cloud App Management server* or step “8” on page 627 of *Upgrading your server*.

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 “14” on page 118 of *Installing your Cloud App Management server* or step “8” on page 627 of *Upgrading your server*). 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 607.

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 “14” on page 118 of *Installing your Cloud App Management server* or step “8” on page 627 of *Upgrading your server*.

## About this task

Use this procedure if you want 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. Change to a different interval, up to 24 hours.

### 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.

When determining the values to set, take into account the considerations in “[Data retention and summarization](#)” on page 607.

## 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
data:
  SHORT_TERM_SUMMARY_TTL: "P60D"
  LONG_TERM_SUMMARY_TTL: "P6M"
  POLICY_GENERATION_HOURLY_INTERVAL: "PT1H"
  DAILY_SUMMARIZATION_ENABLED: "true"
```

where

- *releaseName* is the release name, such as `ibmcloudappgmt`.
  - `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.
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 by issuing 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"
```



# Chapter 15. 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 612.

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

Resource item	Description
	predefined group exists 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>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 597</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 73. 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”</a> on page 606.</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 73. 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 73. 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 Average or a different <b>Aggregation</b> function, such as Maximum or Deviation.</li> <li>3. Select the 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 74. Other monitoring widgets for specific resource types such as Linux or Kubernetes

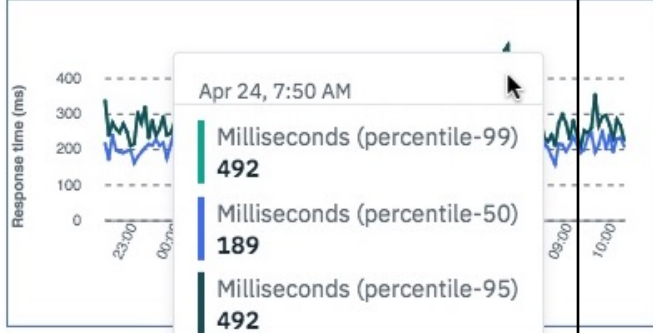

User interface item	Description
<p><b>Line charts and the golden signals</b></p>	<ul style="list-style-type: none"> <li>Line charts plot metrics from the past three hours or as selected. Hover the mouse pointer over a plot point to see the value and time stamp. All the line charts in the dashboard are 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 74. 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 617.</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.

## 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 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 10, “Configuring the ICAM Agents,” on page 161](#) 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.



---

## Chapter 16. 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 V2019.2.1 offering to the IBM Cloud App Management, Advanced V2019.3.0 offering. For more information, see [“Upgrading your server from V2019.2.1 to V2019.3.0”](#) on page 625.

---

### Upgrading your server from V2019.2.1 to V2019.3.0

---

Learn how to upgrade your Cloud App Management server from Advanced V2019.2.1 to Advanced V2019.3.0.

#### Before you begin

When you upgrade your Cloud App Management server, use IBM Cloud Private V3.2.1. The IBM Cloud App Management V2019.3.0 product runs on IBM Cloud Private V3.2.1. For this release, best practice is to upgrade IBM Cloud App Management from V2019.2.1 to V2019.3.0 first, then upgrade IBM Cloud Private from V 3.2.0 to V3.2.1.

**Note:** It is best practice to install the Cloud App Management server in a new, non-default namespace. A `limitrange` resource was 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 did not install the Cloud App Management server in the default namespace, you can disregard this note. If your Cloud App Management server is installed in the default namespace and 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
```

With this release, there are increased system requirements. For more information, see [“System requirements”](#) on page 57

You must download the Cloud App Management installation image file from the IBM Passport Advantage website. Sign in and download the Passport Advantage Archive (PPA) file. For this version, V2019.3.0 the file is named `app_mgmt_server_2019.3.0.tar.gz`. For more information, see [“Part numbers”](#) on page 53.

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 "icamusr" 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.

#### Procedure

Complete the following steps as an IBM Cloud Private cluster administrator:

1. Log in to IBM Cloud Private V3.2.1 and select the namespace where IBM Cloud App Management is installed. Enter 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 for your cluster. The default value is *mycluster*.

Follow the prompts to complete the login.

2. Create an upgrade directory and save or copy the PPA installation file `app_mgmt_server_2019.3.0.tar.gz` from the [IBM Passport Advantage](#) website to the upgrade directory. Change directory to the upgrade directory.

```
mkdir -p upgrade
```

```
cd upgrade
```

3. Extract the Helm charts from the Passport Advantage Archive (PPA) file by running the following commands from the upgrade directory where you saved the `app_mgmt_server_2019.3.0.tar.gz` file:

```
tar -xvf ppa_file charts
tar -xvf charts/ibm-cloud-appmgmt-prod-1.5.0.tgz
```

Where:

- *ppa\_file* Is the compressed Cloud App Management PPA installation image, the `app_mgmt_server_2019.3.0.tar.gz` file.

**Note:** The `charts` value is required and ensures the `tar` command extracts only the `charts` directory from the PPA file. Otherwise, all the images can be extracted leading to potential space issues.

4. Load the PPA file into Docker:

- a) Log in to the Docker private image 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`.

- b) Load the Cloud App Management PPA file into the IBM Cloud Private local repository by running 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 compressed Cloud App Management PPA installation image, the `app_mgmt_server_2019.3.0.tar.gz` file and *my\_helm\_repo\_name* is the name of the target Helm repository. Run the **cloudctl catalog repos** command to get a list of repositories.

**Note:** Unless you specify an `imagePullSecret`, you can access this image from only the namespace that hosts it. If your environment does not use `mycluster.icp:8500` as the registry parameter, you must log in to your registry and specify this parameter. Similarly, you can specify the `--repo` parameter in your environment.

**Note:** This step might take considerable time to complete (>25 minutes).

5. Change directory to the `ibm-cloud-appmgmt-prod` directory:

```
cd ibm-cloud-appmgmt-prod
```

6. Elasticsearch `vm.max_map_count` requirement.

Elasticsearch requires you to set a kernel parameter to run normally. This step needs to be completed on all worker nodes where Cloud App Management is installed. These nodes are determined when you configure the persistent storage. You need to set `vm.max_map_count` to a

value of at least 1048575. Set the parameter with `sysctl` to ensure that the change takes effect immediately:

```
sysctl -w vm.max_map_count=1048575
```

You can also set the parameter in `/etc/sysctl.conf` to preserve the change after a node restart by adding:

```
vm.max_map_count=1048575
```

7. Identify the current release name by entering the following command:

```
helm list --tls | grep ibm-cloud-appmgmt-prod | awk '{print $1}'
```

The default is `ibmcloudappmgmt`. Make a note of the current release name, this value is used for `my_release_name`.

8. Obtain the existing Helm installation's override values and save them to a file named with the current release name:

- a) Enter the following command:

```
helm get values my_release_name --tls > my_release_name-overrides.yaml
```

Where `my_release_name` (for example, `ibmcloudappmgmt`) is the release name from step “7” on page 627.

**Note:** If the `spec.volumeClaimTemplates` values for the `StatefulSet` manifests are edited between any Helm releases, the upgrade fails. This upgrade failure occurs if modifications are made to the release values in the configuration of the persistent volume claims during the upgrade procedure. Do not change the release values that are used in the configuration of the persistent volume claims, such as size of these objects, names .

- b) 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**.

- c) 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.
global:
  metric:
    summary:
      enabled: false
# Metrics summarization is not enabled when set to false.
```

For more information, see “Data retention and summarization” on page 607.

9. Change directory to `PATH/upgrade/ibm-cloud-appmgmt-prod`

Where *PATH* is the directory location where you created the upgrade directory in step “2” on page 626.

10. Create directories for Elasticsearch on all of the workers nodes, create the storage class for Elasticsearch, and create the persistent volumes for all Elasticsearch replicas.

- a) Obtain the persistent volume names for Cassandra. Check to ensure that Elasticsearch is not scheduled on the same worker node as Cassandra to avoid degrading performance:

```
kubectl get pv -lrelease=ibmcloudappmgmt
```

- b) List the YAML files for the Cassandra nodes and get their IP addresses and note them:

```
kubectl get pv ibmcloudappmgmt-cassandra0 ibmcloudappmgmt-cassandra1 ibmcloudappmgmt-cassandra2 -o yaml
```

- c) List all nodes and their IP addresses to ensure the node where Cassandra runs is not used for Elasticsearch, if possible.

```
kubectl get nodes
```

- d) Make a note of the IP addresses of three worker nodes to use for Elasticsearch, then ssh to each node and run the following command to create the elasticsearch directory:

```
mkdir -p /k8s/data/elasticsearch
```

- e) Use ssh to connect to the master node and change directory into the *PATH/upgrade/ibm-cloud-appmgmt-prod* directory.

- f) Create a directory named: *yaml*. Change directory into the *yaml* directory.

```
cd PATH/upgrade/ibm-cloud-appmgmt-prod
mkdir yaml
cd yaml
```

- g) In the *yaml* directory, create YAML files for each worker node you are deploying to. In this case, there are three. When you create the files, you need to ensure that your release-name is the same. In this example case, it is *ibmcloudappmgmt*. Change the files to match your release-name if it is different. Change the metadata name for each node, updating both the release-name, and the numeric value at the end of the string, such as *ibmcloudappmgmt-elasticsearch0*, *ibmcloudappmgmt-elasticsearch1*, *ibmcloudappmgmt-elasticsearch2*, with three Elasticsearch replicas. The IP address will need to be changed to correspond with the node the persistent volume is being deployed to: *PersistentVolume\_icam-elastic\_icam\_metadata\_name.yaml*

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: ibmcloudappmgmt-elasticsearch0
  labels:
    release: ibmcloudappmgmt
spec:
  capacity:
    storage: 50Gi
  storageClassName: ibmcloudappmgmt-local-storage-elasticsearch
  local:
    path: /k8s/data/elasticsearch
  nodeAffinity:
    required:
      nodeSelectorTerms:
        - matchExpressions:
            - key: kubernetes.io/hostname
              operator: In
              values: ["95.95.95.95"]
  accessModes:
    - ReadWriteOnce
  persistentVolumeReclaimPolicy: Retain
```

- h) Create the following file `StorageClass_icam-local-storage-elastic_icam.yaml` in the `ibm-cloud-appmgmt-prod/yaml/` directory. Replace the `release-name` with your release name as necessary:

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: ibmcloudappmgmt-local-storage-elasticsearch
  labels:
    release: ibmcloudappmgmt
provisioner: kubernetes.io/no-provisioner
volumeBindingMode: WaitForFirstConsumer
```

- i) Confirm the StorageClass file and the persistent volume files have been created, and ensure that you are still in the `ibm-cloud-appmgmt-prod/yaml/` directory. Create the persistent volumes for Elasticsearch by running the following command:

```
kubectl apply -f .
```

- j) Change directory into the `PATH/upgrade/ibm-cloud-appmgmt-prod` directory. Edit the `my_release_name-override.yaml` file so helm knows about the new persistent volumes. Make the following changes to the `my_release_name-overrides.yaml` file:

Under the `global.persistence.storageClassOption` key, add the following subkey and value: `elasticdata: ibmcloudappmgmt-local-storage-elasticsearch` replace the `release-name` with your `my_release_name` as necessary. Example:

```
storageClassOption:
  cassandradata: "ibmcloudappmgmt-local-storage-cassandra"
  cassandrabak: "none"
  couchdbdata: "ibmcloudappmgmt-local-storage-couchdb"
  datalayerjobs: "ibmcloudappmgmt-local-storage-datalayer"
  kafkadata: "ibmcloudappmgmt-local-storage-kafka"
  zookeeperdata: "ibmcloudappmgmt-local-storage-zookeeper"
  elasticdata: "ibmcloudappmgmt-local-storage-elasticsearch"
```

Under the `global.persistence.storageSize` key, add the following subkey and value: `elasticdata: 50Gi` Example:

```
storageSize:
  cassandradata: "50Gi"
  cassandrabak: "50Gi"
  couchdbdata: "5Gi"
  datalayerjobs: "5Gi"
  kafkadata: "5Gi"
  zookeeperdata: "1Gi"
  elasticdata: "50Gi"
```

Save and Exit the file `my_release_name-override.yaml`.

11. Run the Helm upgrade to upgrade the existing release by using the `charts/ibm-cloud-appmgmt-prod-1.5.0.tgz` file extracted in step “3” on page 626, the override values file `my_release_name-overrides.yaml` created in step “8” on page 627 and the `my_release_name.2019.3.0.values.yaml` file. Enter the following command:

```
helm upgrade my_release_name charts/ibm-cloud-appmgmt-prod-1.5.0.tgz
--values my_release_name-overrides.yaml --tls
```

Where `my_release_name` is the release name that is obtained in step “7” on page 627, (for example, `ibmcloudappmgmt`).

**Note:** As mentioned in step “8” on page 627 to change Data Retention from the default value of 8 days, you need to add that value back into any of the above override values. Alternatively, you can pass the value in the form: `--set global.metric.retention.rawMaxDays=number_of_days`. For enabling Metrics Summarization, you can edit the override values or you can enable metrics summarization by passing this value in the form: `--set global.metric.summary.enabled=true`. For more information, see “Data retention and summarization” on page 607.

"Happy helming" indicates a successful Helm upgrade. It takes more than ten minutes for all services to finalize their upgrades. Run the `helm status my_release_name --tls` command to monitor the progress.

12. After issuing the **helm upgrade** command, microservices are upgraded. To get the current state of the pods, run the following command:

```
kubectl get pod --namespace my_namespace --selector release=my_release_name
```

Where `my_namespace` is the namespace you selected when logging into IBM Cloud Private.

If you have `watch`, a Linux command line tool, you can use `watch` before the command to reissue it every two seconds by default. During this time, you might be unable to access the Cloud App Management console. If you were already logged in to the Cloud App Management console, you must log out. Then, log in after the new service is ready. There might be a few missing data points from data providers that submitted data to the server during the upgrade time.

## Results

The Cloud App Management server is successfully upgraded.

## What to do next

1. Reconnect your data collectors to the Cloud App Management server following the server upgrade.

For the data collectors, we have the following requirements following a server upgrade:

- The following runtime data collectors require a restart before data is visible on the UI:
    - J2SE data collector
    - Liberty data collector
    - Node.js data collector
    - Python data collector
  - K8s Monitor- these must be updated to 2019.3.0 before data is visible on the UI.
2. Upgrade your Synthetics PoP server. For more information, see [“Upgrading the Synthetics PoP server” on page 640](#).

## Upgrade to a high-availability environment

Upgrade to a high-availability environment.

Before you begin

By default in V2019.2.1, the metric and Jaeger OpenTT data in Cassandra is given a replication factor equal to the number of Cassandra nodes in the cluster, up to a maximum of 3. This changed in 2019.3.0 fresh install to a default maximum of 2, with the option to set to 3 if required. Existing installs are not modified automatically. If you have installed V2019.2.1 with high-availability, modify your replication factor with the following procedure:

1. Identify your Cassandra username and password if needed:

```
kubectl get secret <my_release_name>-cassandra-auth-secret -o json | grep username | awk '{ print $2 }' | sed 's/[",,]//g' | base64 -d
kubectl get secret <my_release_name>-cassandra-auth-secret -o json | grep password | awk '{ print $2 }' | sed 's/[",,]//g' | base64 -d
```

2. Enter the Cassandra container:

```
kubectl exec -it <my_release_name>-cassandra-0 -- /bin/bash
```

3. Run the `cqlsh` shell and enter the decoded password for the user when prompted:

```
/opt/ibm/cassandra/bin/cqlsh -u <my_cassandra_username>
```



#### 4. Alter the replication factor for the 2 keyspaces:

```
ALTER KEYSPACE jaeger_v1_opentt WITH REPLICATION = {'class': 'SimpleStrategy',
'replication_factor': '2'} ;
ALTER KEYSPACE metricdb WITH REPLICATION = {'class': 'SimpleStrategy', 'replication_factor':
'2'} ;
```

#### 5. Verify the replication factor settings:

```
SELECT replication FROM system_schema.keyspaces WHERE keyspace_name='janusgraph';
SELECT replication FROM system_schema.keyspaces WHERE keyspace_name='metricdb';
```

The results should

```
replication
-----
{'class': 'org.apache.cassandra.locator.SimpleStrategy', 'replication_factor': '2'}
```

#### 6. Exit the cqlsh shell by typing exit and enter.

#### 7. While still in the Cassandra container, run repairs against each database:

```
nodetool repair jaeger_v1_opentt
nodetool repair metricdb
```

The upgrade process consists of similar steps to the existing HA install or upgrade paths. Complete the following steps:

1. Upgrade from 2019.2.0 (non-HA) -> 2019.2.1 (non-HA) using the existing normal upgrade procedure described here: [“Upgrading your server from V2019.2.1 to V2019.3.0” on page 625.](#)
2. Upgrade from 2019.2.1 (non-HA) -> 2019.2.1 (HA) by completing the following steps:
  - a. Run `prepare-pv.sh` to prepare new PVs for the additional statefulsets to use. The `prepare-pv.sh` script is described in step [“8” on page 114](#) in the *Installing the Cloud App Management server* topic.

Ensure that you include the existing nodes in the node list, since the script starts numbering at 0. For example, if Cassandra was on node `worker0`, but now you’re adding a `worker1` and `worker2`, run with `--cassandraNodes “worker0 worker1 worker2”`. This replaces the yaml for the PV for `cassandra0`, but it should be identical to the existing PV and therefore it is not modified when the `kubectl create -f ...` command is run in step “9” on page 116 in the Installing the Cloud App Management server topic. For example, if the original install was performed using the following nodes for, putting Cassandra on worker0, and the other PVs on worker1:`

```
./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh --size1 --
cassandraNode "worker0" --zookeeperNode "worker1" --kafkaNode "worker1" --couchdbNode
"worker1" --datalayerNode "worker1"
kubectl create -f /opt/ibm-cloud-private-ee-3.2.0/ibm-cloud-appmgmt-prod/ibm_cloud_pak/
yaml/
```

To move to high availability, 4 new nodes are added, `worker2`, `worker3`, `worker4` and `worker5`. Run the `prepare-pv.sh` command again, including the original `worker0` and `worker1` first in the list. Note, the order is important, as it determines which PV number is assigned, for example:

```
./ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/prepare-pv.sh --size1 --
cassandraNode "worker0 worker2 worker4" --zookeeperNode "worker1 worker3 worker5"
--kafkaNode "worker1 worker3 worker5" --couchdbNode "worker1 worker3 worker5" --
datalayerNode "worker1 worker3 worker5"
kubectl create -f /opt/ibm-cloud-private-ee-3.2.0/ibm-cloud-appmgmt-prod/ibm_cloud_pak/
yaml/
```

#### 3. Complete the following steps to modify the replication factor for Kafka after a Cloud App Management server upgrade.

- a. After adding the additional PVs, upgrade Kafka and its replication factor. First, increase the replica count of Kafka:

```
kubectl scale --replicas=3 statefulset/release_name-kafka
kubectl scale --replicas=3 statefulset/ibmcloudappmgmt-kafka
```

b. Next, enter the Kafka pod in order to execute the replication changes:

```
kubectl exec -it <release_name> -kafka-0 -c release_name -kafka -- /bin/bash
kubectl exec -it ibmcloudappmgmt-kafka-0 -c ibmcloudappmgmt-kafka -- /bin/bash
```

c. Disable additional logging on commands:

```
export KAFKA_LOG4J_OPTS=-Dlog4j.rootLogger=OFF
```

d. Generate a list of all topics, excluding the metric topics, for example:

```
/opt/kafka/bin/kafka-topics.sh --zookeeper $ZOOKEEPER_URL --list 2>/dev/null | egrep
-v "^metric\.json|^metric\.protobuf$" | sed 's/\(.*\)/{ "topic": "\1"} /g; 1 s/^/
{"topics": [ /; $ s/,$/]} /' >/tmp/topics.json
```

e. Generate a new assignment, for example:

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

f. Modify the replication factor by adding brokers to the list of replicas. The following example, is for 3 brokers:

```
sed -i "s/\\"replicas\":"\:[0[^\]]*\]/\\"replicas\":"\:[0,1,2]/g" /tmp/reassignment.json
sed -i "s/\\"replicas\":"\:[1[^\]]*\]/\\"replicas\":"\:[1,2,0]/g" /tmp/reassignment.json
sed -i "s/\\"replicas\":"\:[2[^\]]*\]/\\"replicas\":"\:[2,0,1]/g" /tmp/reassignment.json
```

g. Modify the log\_dirs statements to amend references to “any” to 3 occurrences:

```
sed -i 's/"log_dirs":["any"]}/"log_dirs":["any","any","any"]}' /g' /tmp/reassignment.json
```

h. Run the reassign script, for example:

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

i. Generate a list of metric topics:

```
/opt/kafka/bin/kafka-topics.sh --zookeeper $ZOOKEEPER_URL --list 2>/dev/null | egrep
"^metric\.json|^metric\.protobuf$" | sed 's/\(.*\)/{ "topic": "\1"} /g; 1 s/^/{ "topics":
[ /; $ s/,$/]} /' >/tmp/topics.metric.json
```

j. Generate a new assignment, for example:

```
/opt/kafka/bin/kafka-reassign-partitions.sh --topics-to-move-json-file /tmp/
topics.metric.json --broker-list "0,1,2" --zookeeper $ZOOKEEPER_URL --generate 2>/dev/
null | grep "version" | tail -1> /tmp/reassignment.metric.json
```

k. Modify the replication factor by adding brokers to the list of replicas. For metrics, we are choosing high availability of 2 to reduce the impact of the load on the backend. The following example is for 3 brokers:

```
sed -i "s/\\"replicas\":"\:[0[^\]]*\]/\\"replicas\":"\:[0,1]/g" /tmp/reassignment.metric.json
sed -i "s/\\"replicas\":"\:[1[^\]]*\]/\\"replicas\":"\:[1,2]/g" /tmp/reassignment.metric.json
sed -i "s/\\"replicas\":"\:[2[^\]]*\]/\\"replicas\":"\:[2,0]/g" /tmp/reassignment.metric.json
```

l. Modify the log\_dirs statements to amend references to “any” to 2 occurrences:

```
sed -i 's/"log_dirs":["any"]}/"log_dirs":["any","any"]}' /g' /tmp/reassignment.metric.json
```

m. Run the reassign script, for example:

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

- n. Verify the new replication factors by running the following command (goes through each topic that you modified and prints out a `describe` against it so that you can see if the modification worked):

```
for topic in `cat /tmp/topics.json | grep -o '"topic":.*' | cut -d '"' -f 4` ; do echo  
$topic ; /opt/kafka/bin/kafka-topics.sh --zookeeper $ZOOKEEPER_URL  
--topic $topic --describe 2>/dev/null ; done
```

- o. The output should include ReplicationFactor: 3 (or 2 for metrics), as well as Replicas and Isr (in-sync replicas) lists including all 3 brokers, 0,1,2 (or 2 of 3 for metrics). For example:

```
__consumer_offsets  
Topic: __consumer_offsets      PartitionCount:50      ReplicationFactor:3  
Configs:segment.bytes=104857600,cleanup.policy=compact,compression.type=producer  
Topic: __consumer_offsets      Partition: 0      Leader: 1      Replicas: 1,2,0  
Isr: 0,1,2  
Topic: __consumer_offsets      Partition: 1      Leader: 2      Replicas: 2,0,1  
Isr: 0,1,2  
Topic: __consumer_offsets      Partition: 2      Leader: 0      Replicas: 0,1,2  
Isr: 1,2,0  
Topic: __consumer_offsets      Partition: 3      Leader: 1      Replicas: 1,2,0  
Isr: 0,1,2
```

#### 4. Helm upgrade

- a. After manually modifying Kafka, a helm upgrade needs to be performed to finish scaling up the services. First, generate an override yaml with your existing environment variables

```
helm get values my_release_name --tls > my_release_name-overrides.yaml
```

- b. Preupgrade: preUpgrade.sh can be used to modify the cluster sizes by generating an additional override file:

```
ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/preUpgrade.sh --releaseName  
<my_release_name> --namespace <namespace>3 --kafkaClusterSize 3
```

- c. The output will include the new replica settings, as well as the license acceptance

- d. Use this `my_release_name.2019.2.1.values.yaml` file with the `my_release_nameoverrides.yaml` to run the helm upgrade:

```
helm upgrade --tls ibmcloudappmgmt /opt/ibm-cloud-private-ee-3.2.0/ibm-cloud-appmgmt-  
prod/ --values ibmcloudappmgmt.overrides.yaml --values  
ibmcloudappmgmt.2019.2.1.values.yaml
```

- e. Wait for all of the kafka pods to restart. You can watch their progress with the following command: :  
`watch "kubectl get pods -o wide | grep kafka"`.

5. Now that Kafka is upgraded, the HA upgrade of the remaining services can be performed.

```
helm get values my_release_name --tls > my_release_name-overrides.yaml  
ibm-cloud-appmgmt-prod/ibm_cloud_pak/pak_extensions/preUpgrade.sh --releaseName  
<my_release_name> --namespace <namespace> --minReplicasHPAs 2 --maxReplicasHPAs 3 --  
zookeeperClusterSize 3 --couchdbClusterSize 3 --datalayerClusterSize 3 --  
cassandraClusterSize 3  
helm upgrade --tls ibmcloudappmgmt /opt/ibm-cloud-private-ee-3.2.0/ibm-cloud-appmgmt-prod/ --  
values ibmcloudappmgmt.overrides.yaml --values ibmcloudappmgmt.2019.2.1.values.yaml
```

#### 6. Modify Cassandra replication settings

- a. Identify your Cassandra username and password if needed:

```
kubectl get secret <my_release_name>-cassandra-auth-secret -o json | grep username | awk  
'{ print $2 }' | sed 's/["],//g' | base64 -d  
kubectl get secret <my_release_name>-cassandra-auth-secret -o json | grep password | awk  
'{ print $2 }' | sed 's/["],//g' | base64 -d
```

- b. Enter the Cassandra container:



Management console. For more information, see [“Upgrading your IBM Tivoli Monitoring agents”](#) on page 636.

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 638.

## 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 132.
2. Configure the installation images for communication with the Cloud App Management server. For more information, see [“Configuring the downloaded images”](#) on page 132.
3. Re-install the agent. For more information, see the following topics:
  - [“Installing agents on UNIX systems”](#) on page 134
  - [“Installing agents on Linux systems”](#) on page 138
  - [“Installing agents on Windows systems”](#) on page 141

### 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 162.

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.

–  

```
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, *tf1x8266.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 518 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 532 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 132.
2. Configure the installation images for communication with the Cloud App Management server. For more information, see [“Configuring the downloaded images”](#) on page 132.
3. Re-install the agent. For more information, see the following topics:
  - [“Installing agents on UNIX systems”](#) on page 134
  - [“Installing agents on Linux systems”](#) on page 138
  - [“Installing agents on Windows systems”](#) on page 141

### 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](#). If the data collector dashboards were enhanced for the latest release, data collectors that have not yet been upgraded do not show the enhancements unless the data collectors are restarted. To restart a data collector, and re-register the **Resources** dashboards, use the `kubectl delete pod my_pod` command.

### 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 445.
  - To uninstall the Node.js data collector, see [“Uninstalling the Node.js data collector”](#) on page 455.
  - To uninstall the Liberty data collector, see [“Uninstalling the Liberty data collector from your application”](#) on page 464.
  - To uninstall the J2SE data collector, see [“Uninstalling the J2SE data collector from your application”](#) on page 469.
  - To uninstall the Python data collector, see [“Uninstalling the Python data collector”](#) on page 479.
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 429.
- For instructions about Node.js data collector, see [“Configuring Node.js application monitoring”](#) on page 450.
- For instructions about Liberty data collector, see [“Configuring Liberty application monitoring”](#) on page 456.
- For instructions about J2SE data collector, see [“Configuring J2SE application monitoring”](#) on page 465.
- For instructions about Python data collector, see [“Configuring Python application monitoring”](#) on page 470.

## 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 2019.2.1 or later, 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 611 .
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\_2019.3.0.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 132.
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 17. Troubleshooting (known issues, resolutions, and limitations) 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.

---

### 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 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.

# IBM Multicloud Manager hub cluster fails to import a cluster with Cloud App Management server installed

## Problem

If you have a cluster with the Cloud App Management server installed and you want to import this cluster as a managed cluster into IBM Multicloud Manager, the import fails.

## Cause

On Red Hat OpenShift, a SecurityContextConstraints (SCC) must be bound to the target namespace before you install the Cloud App Management server. The predefined SCC name `ibm-restricted-scc` is deployed (through the CEM subchart) when the `if.Capabilities.APIVersions.Has "security.openshift.io/v1"` capability exists. This SCC has a priority value, which, in this case, is not bound to the IBM Multicloud Manager Klusterlet pods. As a result, the cluster import fails.

## 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: ibmcloudappmgmt-ibm-cem-ibm-restricted-scc
```

## Solution

Edit the `${RELEASE_NAME}-ibm-cem-ibm-restricted-scc` SCC to remove the *Priority* value by setting it to `--validate=false`. The Red Hat OpenShift policy administrator can then assign an SCC to the IBM Multicloud Manager Klusterlet pods. Complete the following steps:

1. Edit the SCC by running the following command:

```
kubectl edit scc ${RELEASE_NAME}-ibm-cem-ibm-restricted-scc --validate=false
```

where `{RELEASE_NAME}` is the name that you gave to your IBM Cloud App Management release.

2. Enter null as the value for the **Priority** parameter.
3. Delete the IBM Cloud App Management SCC by running the following command:

```
kubectl delete scc ibm-cloud-appmgmt-prod-scc -n kube-system
```

## 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**.

## 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.

## Users accounts cannot be edited or saved

---

### Problem

You cannot edit a user for Cloud App Management with IBM Cloud Pak for Multicloud Management. You try to edit, but your edit cannot be saved. The Save button is not enabled.

### Cause

The Save button is not enabled if mandatory data is missing from the account. If user accounts that were synched from IBM Cloud Private are missing mandatory data, the account cannot be edited.

This functions as designed.

## 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 partion 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 `kubectl` (for Kubernetes) and `cqlsh` (for Cassandra) to submit this statement. See the following example:

```
$ kubectl get pods
NAME                                READY   STATUS
cemonicp-cassandra-0                1/1     Running
cemonicp-couchdb-0                  1/1     Running
cemonicp-ibm-cem-brokers-6564b45c47-68pbr 1/1     Running
cemonicp-ibm-cem-users-d4fdf67bc-l2k5b    1/1     Running
cemonicp-ibm-cem-channelservices-65db684f96-x8m9q 1/1     Running
cemonicp-ibm-cem-datalayer-0           1/1     Running
cemonicp-ibm-cem-datalayer-cron-1556759700-xzzz1 0/1     Completed
cemonicp-ibm-cem-event-analytics-ui-7cf786cc99-sqvw8 1/1     Running
cemonicp-ibm-cem-eventpreprocessor-5d8d98dd54-s8nv1 1/1     Running
cemonicp-ibm-cem-incidentprocessor-5bd964646-pkhg6 1/1     Running
cemonicp-ibm-cem-integration-controller-7cb689fffd-kg94x 1/1     Running
cemonicp-ibm-cem-normalizer-74dd497c94-zhlpn 1/1     Running
cemonicp-ibm-cem-notificationprocessor-f9cd8d65d-dql41 1/1     Running
cemonicp-ibm-cem-rba-as-545894d565-s6mn4 1/1     Running
cemonicp-ibm-cem-rba-rbs-5cd7d44556-lp747 0/1     Running
cemonicp-ibm-cem-scheduling-ui-5fbdd6649f-8svn6 1/1     Running
cemonicp-kafka-0                    2/2     Running
cemonicp-redis-sentinel-689c8764b5-9kldm 0/1     CrashLoopBackOff
cemonicp-redis-server-6585cd65cc-277vn 0/1     CrashLoopBackOff
cemonicp-zookeeper-0                1/1     Running
0                                     2m37s
<name>:deploy-ibm-cloud-private <name$> kubectl exec -it cemonicp-cassandra-0 /bin/bash
cassandra@cemonicp-cassandra-0:/$ 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.

## Log Profile Event data not showing correctly

---

### Problem

The Monitoring Agent for Windows OS Log Profile Event data does not show up correctly. The problem is seen with connected Monitoring Agent for Windows OS and configured log file monitoring. The CONF and FMT files are created in the `/localconfig/nt/log_discovery` directory. Check the UI to see matched records and click the log, no data is shown on the Profile Event table. The problem is a known issue.

### Cause

The Monitoring Agent for Windows OS incorrectly determines the value for OSAGNTNODE, which causes Log Profile Event data to not show up correctly.

### Solution

This known issue is being fixed in an interim fix.

## 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.

## Multiple service instances are not supported in same browser

---

### Symptoms

Data in all IBM Cloud App Management web pages are switched to the last service instance registered, when different service instances are logged and accessed in same browser.

### Causes

The browser saves only one instance information, which is the last one registered.

### Environment

Steps to reproduce the behavior:

1. Log in to service instance 1 with a browser A.
2. Log in to service instance 2 in the same browser A.
3. The original service instance 1 is switched to service instance 2 automatically. Take downloading a configuration package from the original service instance 1 as example. It returns information of service instance 2.

### Resolving the problem

Use a separate browser to log in for each service instance. For example:

1. Log in to service instance 1 with a browser A.
2. Log in to service instance 2 with a different browser B.

## No agent metrics are displayed after Kafka restarts

---

### Problem

If Kafka restarts unexpectedly, or is scaled down to 0 and then back up, metric data stops flowing.

### Symptoms

The agentcomm and temacomm containers do not receive data after Kafka is restarted.

### Solution

This problem is a known issue. If Kafka crashes or is restarted, restart the agentcomm and temacomm containers to get the agent data flowing.

## 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.



## 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 `app_mgmt_server_2019.3.0.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.

## How do I reload support files to solve agent or environment problems

You might want to reload updated support files for problems you encounter with agent support files. Alternatively, you might encounter a problem in your environment and you want to reload the existing support files.

### Problem

You might encounter one or more of these issues:

1. You encounter a problem with the agent support files.
2. You encounter a problem in your environment.

### Solution

For the first scenario where you encounter a problem with the agent support files, complete the following steps:

1. Request a new `temasda` docker image from IBM.
2. Then, run the `helm upgrade` to reload the updated support files.

For the second scenario where you encounter a problem in your environment and you want to reload the existing support files, to run the ReLoad API, complete the following steps:

1. Locate the cluster IP address for the temasda service as shown here:

```
kubectl get svc | grep -i temasda
    achantest-temasda ClusterIP 10.0.0.53 <none> 80/TCP,443/TCP
```

2. Use the following **curl** command to run the **reload** command from the Kubernetes master node by using the temasda cluster ip address shown in step 1:

```
curl -i -v -k -X POST "http://10.0.0.53:80/1.0/temasda/reload"
-H "accept: text/plain" -H "content-type: application/json"
```

```
Expected return code: 202
Expected message: Reprocessing service instances from beginning;
may take a few minutes to complete
```

## Results

You ran the helm upgrade to reload the updated agent support files to the server or you used the ReLoad API successfully to reload the existing support files.

## 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.

## 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 123.

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
ibmcloudappmgmt-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 pod_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.

## 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.

## 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.

## 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 unauthorized error

Error messages such as:

```
Error: You must be logged in to the server (Unauthorized)
Error: Unauthorized
```

## Symptoms

This error occurs 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
cloudctl cluster-config mycluster
```

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
```

## Error during console login

IBM Cloud Private uses an Open ID Connect (OIDC) authorization provider to validate user credentials during login. If, after entering your login credentials to start the Cloud App Management console, you get



an error message about the provider not getting the origin token, you might need to perform OIDC registration manually.

## Problem

The pods might be unable to communicate externally to the Master IP, causing the OIDC registration to fail.

## Symptom

After entering your user credentials to log in to the Cloud App Management console, you receive the following error:

```
{"code":401,"message":"CEM auth provider middleman could not get origin token.",  
"level":"error"}
```

## Solution

Follow these steps to run the OIDC registration manually:

1. Run the following command to retrieve the OIDC payload json file:

```
kubectl exec -n my_namespace -t `kubectl get pods -l release=my_releasename  
-n my_namespace | grep "my_releasename-ibm-cem-cem-users" | grep "Running" | head -n 1 | awk  
'{print $1}'`  
cat -- "/etc/oidc/oidcPayload.json" > /tmp/oidcPayloadTemplate.json
```

where *my\_namespace* and *my\_releasename* are the namespace and release name that were installed within IBM Cloud Private.

2. Retrieve the *oidcclientid* and *oidcclientsecret*:

```
echo $(kubectl get secret my_releasename-cem-cemusers-cred-secret -o yaml -n  
my_namespace | grep oidcclientid: | awk '{print $2}') | base64 --decode  
echo $(kubectl get secret my_releasename-cem-cemusers-cred-secret -o yaml -n  
my_namespace | grep oidcclientsecret: | awk '{print $2}') | base64 --decode
```

where *my\_releasename* and *my\_namespace* are the release name and namespace installed within IBM Cloud Private.

3. Update the */tmp/oidcPayload.json* with the *clientid* and *clientsecret*:

```
sed -e "s/\\${client_id}/my_oidcclientid/" -e "s/\\${client_secret}/  
my_oidcclientsecret/" /tmp/oidcPayloadTemplate.json > /tmp/oidcPayload.json
```

where *my\_oidcclientid* and *my\_oidcclientsecret* are from the previous step.

4. As a cluster administrator, retrieve the *oauthadmin* password:

```
echo $(kubectl get secret platform-oidc-credentials -o yaml -n kube-system | grep  
OAUTH2_CLIENT_REGISTRATION_SECRET: | awk '{print $2}') | base64 --decode
```

5. As a cluster administrator, register or update the *oidcclientid*:

```
curl -k -X POST -u oauthadmin:my_oauthadmin-password -H "Content-Type:  
application/json" --data @/tmp/oidcPayload.json https://my_masterip:9443  
/oidc/endpoint/OP/registration
```

```
curl -k -X PUT -u oauthadmin:my_oauthadmin-password -H "Content-Type:  
application/json" --data @/tmp/oidcPayload.json https://my_masterip:9443  
/oidc/endpoint/OP/registration/my_oidcclientid
```

where *my\_oauthadmin\_password* and *my\_oidcclientid* are from the previous steps, and *my\_masterip* is the cluster Master IP.

## Results

After you enter the **curl POST** or **PUT** command to register or update the *my\_oidcclientid*, the following response is provided and you can now log in:

```
root@fvt2ldap:~/vlaunch-scao-fvt# kubectl exec -n scao-fvt -t `kubectl get pods
-l release=scao -n scao-fvt | grep "scao-ibm-cem-cem-users" | grep "Running"
| head -n 1 | awk '{print $1}'` bash -- "/etc/oidc/oidc_reg.sh"
" `echo $(kubectl 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    1091    100    1091     0       0    2029       0  --:--:--  --:--:--  --:--:--   2027

✓ Client exists...
Updating registration...
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
   100    2076    100    1105    100     971    5771    5071  --:--:--  --:--:--  --:--:--   5785

Done.
```

For more information about user authentication and OIDC, see the *Authentication* topic under [Managing access to your platform](#) in the IBM Cloud Private topic collection for IBM Knowledge Center.

## 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.

## 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, *ibmcloudapppgmt* 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.

## V6 WebSphere Applications agent can't start after applying IF0004 agent patch on AIX 7.1.

---

### Problem

The WebSphere Applications agent can't start after applying IF0004 agent patch on AIX 7.1.

## Symptom

### Error description in agent log file:

```
(5B55B651.0001-1:kwjvm.cpp,555,"KwjJvm::hasException") ERROR: caught Java exception
(5B55B651.0002-1:kwjvm.cpp,574,"KwjJvm::hasException") java.lang.UnsatisfiedLinkError:
  net (Not found in com.ibm.oti.vm.bootstrap.library.path)
(5B55B651.0003-1:kwjagent.cpp,310,"KwjAgent::start") ERROR: internal exception in src/bridge/
  jni/kwjjni.cpp at line 112
(5B55B651.0004-1:kwjagent.cpp,312,"KwjAgent::start") Agent stopped
```

## Cause

JVM fails to maintain the LIBPATH updates that are made during the JVM initialization.

## Solution

1. Edit `/opt/IBM/ITM/config/yn.ini`.
2. Add the IBM Tivoli Monitoring JRE library path into LIBPATH property, for example:

```
ITM_Home/JRE/platform_code/lib/ppc:ITM_Home/JRE/platform_code/lib/ppc/j9vm
```

Where

ITM\_Home is the agent installation location

platform\_code is current OS-specific code.

**Note:** The V6 WebSphere Applications agent is 32-bit and uses 32-bit JRE. For example, if the 32-bit JRE is under the aix523 directory, then the IBM Tivoli Monitoring JRE library path is:

```
/opt/IBM/ITM/JRE/aix523/lib/ppc:/opt/IBM/ITM/JRE/aix523/lib/ppc/j9vm
```

## 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.

## Fresh installed Cloud APM Db2 agent on Windows fails to collect data

---

### Problem

Fresh installed Cloud APM Db2 agent on Windows fails to collect data.

### Symptom

The agent logs show the following messages:

```
(5BC44BBD.0002-1:kg1cry.c,2225,"readKeyFile") Cannot read $(CANDLEHOME)/keyfiles\KAES256.ser
(5BC44BBD.0003-1:kg1cry.c,3507,"CRY_Decrypt") Function failed with error code 24
...
'SQL30082N Security processing failed with reason "24" ("USERNAME AND/OR
+5BCF7D4E.002C PASSWORD INVALID"). SQLSTATE=08001'
```

Agent logs is located at *CANDLE\_HOME*\logs. *CANDLE\_HOME* is the agent installation directory.

## Solution

1. Locate the instance specific environment file `CANDLE_HOME\TMAITM6_x64\KUDENV_InstName`.

Where:

`CANDLE_HOME` is the agent installation directory.

`InstName` is the agent instance name.

2. Replace all occurrences of `${CANDLEHOME}` with the actual agent installation directory path.

Example of actual agent installation directory path on Windows is `C:\IBM\APM`

3. Restart the agent

**Note:** Repeat the above steps for each of the existing agent instances where **username** and **password** were provided during configuration.

## Permission error on the `ud.env` file when starting the second instance of APM Db2 agent with TEMA patch

---

### Problem

There is permission error on the `ud.env` file when starting the second instance of APM Db2 agent with TEMA patch.

### Symptom

The terminal shows the following message when starting the agent:

```
/opt/ibm/apm/agent/tmp/ud_db2apm.run: line 226: /opt/ibm/apm/agent/logs/ud.env: Permission denied
```

### Solution

After starting the first instance of the agent, the user must run `CANDLE_HOME/bin/secure.sh` to grant permissions to the admin group that configured the DB2 agent instances. `CANDLE_HOME` is the agent installation directory.

1. Run the `secure.sh` command with the following parameters:

```
./secure.sh -g Db2_admin_group_name
```

Where:

- `-g` indicates group.
- `Db2_admin_group_name` is the Db2 admin group name.

For example: `./secure.sh -g db2iadm1`

2. Start the second instance of 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.

## 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 597](#).

## Dashboards for Kubernetes Resources are empty after Cloud App Management server upgrade

---

## Symptom

After you update the Cloud App Management server 2019.2.0 to 2019.2.1, the dashboards for such runtime data collectors as the Liberty data collector and Node.js data collector show no data and the UI pod log shows errors connecting to Couch: `kubectl logs {amui pod} -c amuirest-service`. You try running the **kubectl delete pod *my\_pod\_name*** command, which restarts the runtime data collector to re-register the **Resources** dashboards, but the dashboards are still empty.

## Cause

Connectivity with the CouchDB was lost during 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 showing no metrics, restart the UI pod: `kubectl delete pod amui`

## 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.

## 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.

## 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`.



If you are not actively working in your

## 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.

## 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 timestamp 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 get help to troubleshoot issues, you can contact IBM Support to open a support ticket.

### About this task

You can open a new service request easily to help you to troubleshoot issues when they arise.

### Procedure

1. Go to the [IBM Support](#) site.
2. In the **Problem Ticketing (PMRs)** section, click **Service Request**.
3. Create or sign in with your IBMId and create your service request.

---

## Chapter 18. 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 19. 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 20. 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. Select a node . js test script. For information about creating a node . js script to test, see [“Create a REST API test case”](#) on page 487.

### Response validation

4. Configure the warning and critical events conditions for your synthetic test in the **Response Validation** section. Accept the defaults, or edit the **Value** and **Unit** for each row. Response times that exceed 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 the Alert Triggers step.

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 496.

### 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 **Private 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 Private Location, see [“Installing a Synthetics PoP”](#) on page 480.

### 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 **Failure detected** to **Off**. To increase the number of failures allowed before a critical alert is triggered, select **consecutively** under **Failure detected**. The default number of consecutive failures is 2, but this number can be customized.

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, select **consecutively** under **Threshold breached**. The default number of slow response times is 2. This number can be customized.

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 496](#).



# 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](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>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>.





SC27950505

