

IBM System Storage N series



Virtual Storage Console 4.1 for VMware vSphere Installation and Administration Guide

Contents

Preface	1
Supported features	1
Websites	1
Getting information, help, and service	1
Before you call.	1
Using the documentation	2
Hardware service and support	2
Firmware updates	2
How to send your comments	2
Virtual Storage Console for VMware vSphere Overview	3
Interaction between the capabilities.	4
Architecture of Virtual Storage Console for VMware vSphere	4
Methods for accessing the capabilities	5
VSC for VMware vSphere provides lock management for system resources	6
Online help.	6
Installation overview	7
VSC for VMware vSphere supported configurations	8
Installing VSC for VMware vSphere using the installation wizard	9
Installing VSC for VMware vSphere using silent mode	10
Registering VSC for VMware vSphere with vCenter Server	11
VSC for VMware vSphere port requirements	11
Regenerating an SSL certificate for VSC for VMware vSphere.	12
How to configure role-based access control for VSC for VMware vSphere	13
Preparation required before upgrading VSC for VMware vSphere	14
Upgrading VSC for VMware vSphere.	15
Adding or removing the Backup and Recovery capability	16
Uninstalling VSC for VMware vSphere using Add/Remove Programs.	17
Uninstalling VSC for VMware vSphere using silent mode	17
Monitoring and configuring hosts and storage controllers	19
VSC for VMware vSphere configuration	20
Storage system discovery and credentials overview	20
Default controller credentials simplify administration of capabilities	21
Specifying credentials with Monitoring and Host Configuration	21
Enabling discovery and management of vFiler units.	22
Enabling discovery and management of vFiler units on private networks.	23
Tunneled vFiler units and Vservers discovered automatically.	23
Differences between direct connections to Vservers and to cluster-management LIFs	23
Discovering and adding storage resources	24
Correcting controller names displayed as "unknown"	25
Removing controllers from the Monitoring and Host Configuration capability	25
Administering the Monitoring and Host Configuration capability	26
Inventory panel selection limits what VSC for VMware vSphere displays.	26
Configuring ESX server multipathing and timeout settings	27
ESX host settings set by Monitoring and Host Configuration capability	27
UNMAP setting turned off in ESX 5.x	28
Direct path access and NFS datastores	29
Changing NFS data paths to direct access	29
NFS Plug-in for VMware VAAI requires additional installation steps	30
Using Web-based tools to manage storage	31
The Monitoring and Host Configuration capability displays configuration details	32
MultiStore vFiler units are displayed differently	32
Collecting diagnostic information	33

Changing the service account for data collection on Windows Server 2008	33
Downloading and running tools	34
Enabling the ESXi secure shell	34
Downloading and installing MBR tools for ESXi hosts	35
Downloading and installing MBR tools for ESX hosts	36
Installing GOS scripts	37
Fields and commands described by the online help	38
Overview panel fields and commands	38
Storage Details - SAN panel fields and commands	47
Storage Details - NAS panel fields and commands	48
Data Collection panel fields and commands	51
Tools panel fields and commands	52
Discovery Status panel fields	53
Provisioning and cloning datastores and virtual machines	57
Tips for working with Provisioning and Cloning	57
Cloning and managing virtual machines.	57
Cloning virtual machines	58
Managing connection brokers	63
Adding connection brokers	63
Removing connection brokers	64
Redeploying clones (locally)	64
Removing a baseline	65
Reclaiming space on virtual machines	65
Importing virtual machines into XenDesktop	66
Importing the file into XenDesktop 4	66
Importing the file into XenDesktop 5	67
Managing storage controllers	67
(Data ONTAP operating in 7-Mode) Viewing storage controller details	67
(Data ONTAP operating in 7-Mode) Removing or adding network interfaces, volumes, and aggregates.	67
Managing volume settings	68
Managing datastores	69
(Data ONTAP operating in 7-Mode) Replicating datastores to remote sites	69
Setup Replication	70
Remove a datastore replication relationship.	71
Synchronize	71
Define or modify a SnapMirror schedule	71
Refreshing the display.	72
Provisioning datastores	72
Mounting datastores	73
Managing deduplication	74
Resizing datastores	74
Destroying datastores	75
About multiple datastores (NFS only)	75
Provisioning and Cloning support files	75
Preferences File	75
Provisioning and Cloning logs	81
Log configuration file	82
Modifying logging levels	82
Using a different log configuration file	82
Export Files	82
Network configuration file	82
XenDesktop export file	83
Provisioning and Cloning vCenter privileges	84
Optimizing and migrating datastores and virtual machines	87
Types of alignments	88
Important notes about using the Optimization and Migration capability	89
The Optimization and Migration workflow	91
Scanning the datastores	91

Scheduling a scan of datastores	92
Performing an online alignment	93
Migrating virtual machines	94
Backing up and restoring data	97
Backup and Recovery requirements	97
Backup and Recovery requirements for optional SnapMirror protection	97
Configuring the Backup and Recovery capability	98
Authentication methods in the Backup and Recovery capability	98
Custom user accounts for accessing a storage system	98
Creating custom users	99
Creating a custom storage system user account	99
How the Backup and Recovery capability discovers vFiler units	99
Managing backups	100
Considerations for adding a backup job	100
Backing up a virtual machine	100
Backing up a datastore or datacenter	101
Starting a one-time backup	103
Editing a backup job	103
Deleting a scheduled backup job	103
Suspending an active backup job	104
Resuming a suspended backup job	104
Restoring data from backups	104
Where to restore a backup	105
Restore operations using data that was backed up with failed VMware consistency snapshots	105
Restoring data from backups	105
Searching for backups	106
Restoring a datastore	106
Restoring a virtual machine or its VMDKs	106
Mounting a backup	107
Unmounting a backup	107
Single file restore	108
How Virtual Storage Console detects network connectivity	108
The difference between limited and direct connectivity	108
Types of file restore sessions	109
Manually creating a .sfr file for the Restore Agent	109
General configuration settings for single file restore	110
Setting session defaults	110
Changing the network connection for a port group	110
Setting the SnapManager for Virtual Infrastructure server address	110
Self-service example workflow	110
Creating a self-service restore session	111
Installing Restore Agent	112
Load the configuration file	112
Recovering single files from a virtual machine	112
Clear the configuration	113
Limited self-service example workflow	113
Create a limited self-service restore session	113
Installing Restore Agent	114
Configuring vCenter Server	115
Recovering single files from a virtual machine	115
Clear the configuration	116
VSC CLI commands	116
Launching the VSC CLI	116
smvi backup create	117
smvi backup delete	118
smvi backup list	119
smvi backup mount	120
smvi backup rename	121
smvi backup restore	122
smvi backup unmount	123

smvi discover datastores	124
smvi filerestore add-portgroup	125
smvi filerestore delete-portgroup	126
smvi notification list	126
smvi notification set	127
smvi notification test	127
smvi restoreagent set	128
smvi servercredential delete	128
smvi servercredential list	129
smvi servercredential set	129
smvi storagesystem add	130
smvi storagesystem delete	130
smvi storagesystem list	131
smvi storagesystem modify	132
smvi version	132
Programmable APIs	135
What the programmable APIs are	135
What you can do with the APIs for VMware vCloud	135
Provisioning and Cloning programmable API	135
The virtual machine clone engine.	135
The datastore management engine	136
The file copy/clone offload engine	136
Provisioning and Cloning methods	136
Virtual machine clone creation and redeploy engine	136
Datastore management engine.	139
Connection Broker features.	142
Copy/Clone offload engine.	144
Utility methods.	149
Provisioning and Cloning specifications and messages.	152
RequestSpec.	153
CloneSpec	153
VmFileSpec	156
DatastoreSpec	156
ControllerSpec	159
VmSpec	160
GuestCustomizationSpec	161
ConnectionBrokerSpec	162
StatusMessage	163
Provisioning and Cloning sample code	163
Provisioning and Cloning client-side programming.	164
Accessing the SOAP API through Java	164
Accessing SOAP through C#	165
Troubleshooting	167
Issues that apply to multiple capabilities	167
Check the Release Notes.	167
VMware only supports selecting one object when using right-click actions	167
Issues that apply to the Monitoring and Host Configuration capability	167
Getting information about storage controllers with an Alert status.	167
Getting information about an ESX and ESXi host with an Alert status	168
Collecting the VSC for VMware vSphere log files	169
Troubleshooting error message "The client cannot communicate with the Virtual Storage Console Server"	169
Updating vCenter credentials for background discovery	170
Resolution of issues with the Backup and Recovery capability	170
Backup and Recovery capability values that you can override	170
Backup and Recovery capability event and error logs	171
Email notification for scheduled backup contains a broken link.	171
You may have reached the maximum number of NFS volumes configured in the vCenter.	171

Using ESX hosts with IBM N series storage	173
LUN type guidelines	173
Manually provisioning storage	173
How to set up VMware ESX	174
Configuring the VMware ESX host	174
(Data ONTAP operating in 7-Mode) Manually setting the path selection policy for Microsoft cluster configurations	175
Timeout values for guest operating systems	175
Running the GOS timeout scripts for Linux	176
Running the GOS timeout scripts for Solaris	177
Running the GOS timeout script for Windows	177
How to identify and fix VMDK partition alignment issues	178
Checking VMDK partition alignment with mbralign	178
VMDK partition alignment with mbralign overview	179
Fixing VMDK partition alignment using mbralign	181
Reinstalling GRUB for Linux guests after running mbralign	182
 Websites	 185
 Copyright and trademark information	 187
Trademark information	188
 Notices	 189
 Index	 191

Preface

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website (accessed and navigated as described in Websites).

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:
www.ibm.com/storage/nas/
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:
www.ibm.com/storage/support/nseries/
This web page also provides links to AutoSupport information as well as other important N series product resources.
- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:
www.ibm.com/systems/storage/network/interophome.html
- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:
publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.

- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in Websites) for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in Websites).

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in Websites).

Note: If you do not see new firmware updates on the N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by email to starpubs@us.ibm.com.

Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

Virtual Storage Console for VMware vSphere Overview

Virtual Storage Console for VMware vSphere software is a single vCenter Server plug-in that provides end-to-end lifecycle management for virtual machines in VMware environments using IBM N series storage.

The plug-in provides different capabilities to perform the following functions:

- **Storage configuration and monitoring using the Monitoring and Host Configuration capability**

The Monitoring and Host Configuration capability handles adding and removing storage controllers and assigning storage controller credentials for all the capabilities except Backup and Recovery. It also enables you to manage ESX and ESXi servers connected to IBM N series storage. You can set values for host timeouts, NAS, and multipathing as well as view storage details and collect diagnostic information.

- **Datastore provisioning and virtual machine cloning using the Provisioning and Cloning capability**

The Provisioning and Cloning capability uses FlexClone technology to let you efficiently create, deploy, and manage the lifecycle of virtual machines from an interface that has been integrated into the VMware environment.

- **Online alignments and single and group migrations of virtual machines into new or existing datastores using the Optimization and Migration capability**

The Optimization and Migration capability enables you to quickly check the alignment status of virtual machines. If there are alignment issues with virtual machines in VMFS datastores, you can in most cases resolve those issues without having to power down the virtual machines.

- **Backup and recovery of virtual machines and datastores using the Backup and Recovery capability**

The Backup and Recovery capability allows you to rapidly back up and recover multihost configurations on IBM N series storage.

VSC for VMware vSphere also provides APIs for VMware vCloud, which enable access to its capabilities in the context of vCloud Director objects and semantics. These APIs enable you to manage credentials for multiple vCenter Servers, discover vCloud Director objects for vCloud tenants, and provision and clone vApps.

VSC for VMware vSphere also supports the Provisioning and Cloning API, which is designed to be leveraged with the VI SDK.

As a vCenter Server plug-in, VSC for VMware vSphere is available to all vSphere Clients that connect to the vCenter Server. Unlike a client-side plug-in that must be installed on every vSphere Client, you install the VSC for VMware vSphere software on a Windows server in your data center.

Note: Do not install this software on a client computer.

The software adds an IBM N series icon to the Solutions and Applications panel of the vSphere Client home page.

When you select the About panel, VSC for VMware vSphere displays its version information as well as version information for each of the installed capabilities.

Related concepts:

“Monitoring and configuring hosts and storage controllers” on page 19

“Provisioning and cloning datastores and virtual machines” on page 57

“Optimizing and migrating datastores and virtual machines” on page 87

“Backing up and restoring data” on page 97

Interaction between the capabilities

All of the VSC for VMware vSphere capabilities provide functions you can use to manage your VMware environment. In most cases, the VSC for VMware vSphere capabilities operate separately from each other. There are some areas, though, where the capabilities interact with each other.

These areas include the following:

- There is a single installation program that you use to install the capabilities. The installation program automatically installs the Monitoring and Host Configuration, Provisioning and Cloning and Optimization and Migration capabilities. If you want to install the Backup and Recovery capability, you must select it when the installation program starts.
- The Monitoring and Host Configuration capability manages storage system discovery and removal and the default storage controller credentials for the Provisioning and Cloning and Optimization and Migration capabilities. The Backup and Recovery capability discovers its own set of storage systems and maintains its own list of credentials.
- VSC for VMware vSphere provides lock management for the capabilities, so as to prevent two capabilities from acting on the same VM or datastore at the same time. As a result, some alignment, migration, provisioning, cloning, and recovery features become unavailable when multiple capabilities attempt to use the same target VM or datastore at the same time.
- Both the Monitoring and Host Configuration capability and the Backup and Recovery capability support vFiler tunneling on physical storage systems that contain vFiler units.

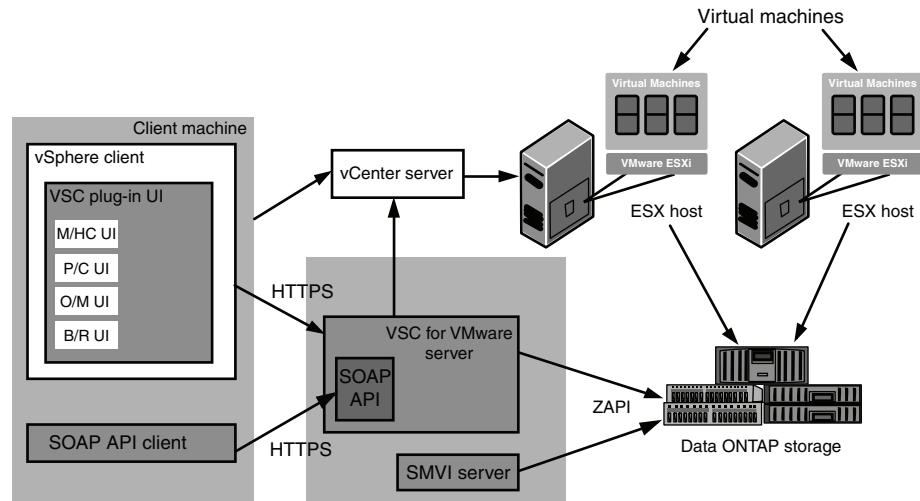
Architecture of Virtual Storage Console for VMware vSphere

The VSC for VMware vSphere architecture includes the storage system running Data ONTAP, the vCenter Server, the vCenter client, the ESX or ESXi host, and the VSC for VMware vSphere capabilities.

VSC for VMware vSphere uses VMware-recommended web-based architecture. It consists of two major components:

- A graphical user interface (GUI) web application that displays as a plug-in within the vSphere client, thus providing a single management console for virtualized environments.
- A server component that is controlled by the VSC for VMware vSphere service and hosts Java Servlets to handle the GUI and API calls to and from the storage systems and the ESX/ESXi hosts.

The following diagram provides a high-level overview of the VSC for VMware vSphere architecture:



When you run VSC for VMware vSphere, you use the VMware vSphere client and the VMware vCenter server. VSC for VMware vSphere provides the following:

- Plug-ins for the four capabilities: Monitoring and Host Configuration, Provisioning and Cloning, Optimization and Migration, and Backup and Recovery. Each capability has its own user interface and online help.
- The VSC for VMware vSphere server
- The SMVI server

In addition you can write applications that communicate with the VSC for VMware vSphere server. For example, you can use the SOAP API to create a client. VSC for VMware vSphere supports the VSC for VMware vSphere API, which works in a VMware vCloud Director deployment and the Provisioning and Cloning API.

The vSphere client and any applications you create use the HTTPS protocol to communicate. The VSC for VMware vSphere server and the SMVI server use ZAPI to communicate with the storage systems that are running Data ONTAP. Communication from the VSC for VMware vSphere server to the vCenter is done using SOAP.

The vCenter server communicates with the physical servers where ESX or ESXi hosts are running. You can have multiple virtual machines running on the ESX or ESXi hosts. Each virtual machine can run an operating system and applications. The ESX and ESXi hosts then communicate with the storage systems.

Methods for accessing the capabilities

VSC for VMware vSphere provides a graphical user interface (GUI) that lets you move between the capabilities so that you can always access the capability you need.

You can click the IBM N series icon in the vCenter Server and then choose the capability you want from the navigation pane.

If you know the task you want to perform, you can right-click on an object in the Inventory panel and then select the IBM N series tab. From that tab, you can select the action you want to perform.

VSC for VMware vSphere provides lock management for system resources

VSC for VMware vSphere uses lock management to keep multiple capabilities from performing simultaneous tasks on the same target datastores or virtual machines. As a result, certain alignment, migration, provisioning, cloning, and recovery features become temporarily unavailable if the target datastore or virtual machine is already being used by another capability.

The lock management process can be seen when you are working with the Provisioning and Cloning capability, the Optimization and Migration capability, and the Backup and Recovery capability. For example, if you are migrating VMs using the Optimization and Migration capability, you cannot use the Provisioning and Cloning capability to clone one of those VMs until the migration is complete.

You cannot perform the following Backup and Recovery tasks when the target datastore or VM is in use by the Provisioning and Cloning capability or the Optimization and Migration capability:

- Recover a datastore, VM, or virtual disk file
- Mount a backup for file restore session
- Unmount a backup that was previously mounted for a file restore session

Note: When a lock occurs on a file restore mount or unmount operation, the lock is held from the time the backup is mounted to the virtual machine until the time the backup is unmounted.

Conversely, if the Backup and Recovery capability has already begun one of the operations listed above, the datastore or VM is unavailable for Provisioning and Cloning operations and Optimization and Migration operations.

Related concepts:

“Provisioning and cloning datastores and virtual machines” on page 57

“Backing up and restoring data” on page 97

Online help

The VSC for VMware vSphere GUI has a separate online Help module for each capability. Each online Help module describes the fields and commands for its capability.

You can access online Help for all the VSC for VMware vSphere capabilities from the vSphere Client **Help** menu:

- **Help > IBM > Monitoring and Host Configuration > Monitoring and Host Configuration Help**
- **Help > IBM > Provisioning and Cloning > Provisioning and Cloning Help**
- **Help > IBM > Optimization and Migration > Optimization and Migration Help**
- **Help > IBM > Backup and Recovery > Backup and Recovery Help**

The Help information is displayed in a web browser.

Installation overview

You can install VSC for VMware vSphere on a 32-bit or 64-bit Windows server.

Note: Do not install this software on a client computer.

To install the VSC for VMware vSphere software, follow these guidelines:

- VSC for VMware vSphere must be installed on a local disk of the Windows server. Do not attempt to install VSC for VMware vSphere on a network share.
- VSC for VMware vSphere requires network connectivity between all of the components it manages.
- A reboot is not required to complete the installation. However, vSphere Clients must be closed and restarted to be able to display the VSC for VMware vSphere Plug-in.
- At a minimum, the display must be set to 1,280 by 1,024 pixels to view VSC for VMware vSphere pages correctly.

By default, VSC for VMware vSphere installs the following three capabilities:

- Monitoring and Host Configuration
- Provisioning and Cloning
- Optimization and Migration

You have the option of also installing the Backup and Recovery capability. If you choose the Backup and Recovery capability, you must purchase a license for SnapManager for Virtual Infrastructure.

The following software licenses might be required for VSC for VMware vSphere depending on which capabilities you use:

- The required protocol license (NFS, FCP, iSCSI)
- SnapManager for Virtual Infrastructure (if installing the Backup and Recovery capability)
- SnapMirror (if using the Provisioning and Cloning capability or if using the SnapMirror update option in the Backup and Recovery capability)
- SnapRestore (if installing the Backup and Recovery capability)
- A_SIS (if using the Provisioning and Cloning capability when configuring deduplication settings)
- MultiStore (if using the Provisioning and Cloning capability and working with vFiler units)
- FlexClone

The FlexClone license is required in the following situations:

- You are using the Provisioning and Cloning capability to clone virtual machines.
- You are using the Backup and Recovery capability in NFS environments, and running a version of Data ONTAP prior to 8.1.

You do not need a FlexClone license if you are running the Backup and Recovery capability in NFS environments with one of the following versions of Data ONTAP 8.1:

- Data ONTAP 8.1 operating in 7-Mode
- Data ONTAP 8.1.1 operating in Cluster-Mode

VSC for VMware vSphere supported configurations

VSC for VMware vSphere is supported on specific releases of ESX/ESXi and Data ONTAP software.

Server configuration

Your Windows system must meet minimum hardware requirements before installing the VSC for VMware vSphere software.

The memory requirements depend on whether you install VSC for VMware vSphere on the same machine as the vCenter Server or on a different machine. When this document was created, the memory requirements for 32-bit environments and 64-bit environments where VSC for VMware vSphere was installed on a separate machine were the following:

- Minimum memory requirement: 4 GB RAM
- Recommended memory requirement: 4 GB RAM

Hardware requirements are higher if you are running VSC for VMware vSphere on the same machine as the vCenter Server.

Note: Please refer to VMware documentation for the current list of hardware requirements.

You should be aware of the following requirements before you install the VSC for VMware vSphere software:

- Supported Microsoft Windows software
- vCenter Server requirements
- ESX host software requirements
- Data ONTAP requirements

See the Interoperability Matrix, which is online at www.ibm.com/systems/storage/network/interophome.html, for details.

Note: VSC for VMware vSphere and the Backup and Recovery Restore Agent do not support IPv6. If the server on which you are installing VSC for VMware vSphere has IPv6 enabled, IPv6 should be disabled before installing VSC for VMware vSphere.

If you are using the single file restore (SFR) feature, you might also have to change a configuration setting to prevent the SFR restore session email from sending an IPv6 address. See the Known issues: Backup and Recovery section of the *Virtual Storage Console for VMware vSphere Release Notes* for the configuration workaround.

Client configuration

The client computer that runs the VMware vSphere Client software must have Microsoft Internet Explorer 8 installed.

Installing VSC for VMware vSphere using the installation wizard

You can use the installation wizard to install VSC for VMware vSphere. By default, the VSC for VMware vSphere software installs the Monitoring and Host Configuration, Optimization and Migration, and Provisioning and Cloning capabilities and gives you the option of installing the Backup and Recovery capability.

Before you begin

You must be logged on with administrator privileges to the machine where you are installing VSC for VMware vSphere. If you attempt to register VSC for VMware vSphere without having administrator privileges, the task will not complete correctly.

You can perform an upgrade if you are running version 2.0 or later of VSC for VMware vSphere. The VSC for VMware vSphere installer does not support upgrades from versions of VSC for VMware vSphere prior to 2.0 or from standalone versions of the Rapid Cloning Utility (RCU), which is now called Provisioning and Cloning, or SMVI (now called Backup and Recovery). If you have any of that software installed, you must uninstall it before you install the current version of VSC for VMware vSphere. If the VSC for VMware vSphere installer finds one of those versions of VSC for VMware vSphere, RCU, or SMVI on the server, it prompts you to uninstall the software, and then aborts.

Note: Make sure you are using the correct installer package for your host machine. You cannot run the 32-bit installer on a 64-bit host machine, or the 64-bit installer on a 32-bit host machine.

Procedure

1. Obtain the product software by downloading the software as follows:
 - a. Go to the IBM NAS support Web site.
 - b. Sign in with your IBM ID and password. If you do not have an IBM ID or password, or if you are adding new N series machines and serial numbers to an existing registration, click the **Register** link, follow the online instructions, and then sign in.
 - c. Select the N series software you want to download, and then select the **Download** view.
 - d. Use the **Software Packages** link on the Web page presented and follow the online instructions to download the software to a working directory on the Windows server.
2. Double-click the installer icon, and click **Run** to start the installation wizard.
3. Follow the instructions in the installation wizard to install the software.

Note: If you want to install the Backup and Recovery capability, you must select that option. Otherwise, the VSC for VMware vSphere installer does not install that capability.

4. Click **Finish** to complete the installation. A Web page appears when the installation is complete. You must register VSC for VMware vSphere with the vCenter Server. You must provide the vCenter Server host name or IP address and the administrative credentials.

Note: If you are not logged on with administrator privileges, you might encounter a problem during the registration process.

What to do next

If you uninstall SMVI prior to installing VSC for VMware vSphere, then before you migrate the backup metadata, you must save this metadata by copying the contents installed in the C:\Program Files\IBM\SMVI\server\repository folder into the C:\Program Files\IBM\Virtual Storage Console\smvi\server\repository folder.

After you finish copying the backup metadata, you must copy the contents of the credential file from the C:\Program Files\IBM\SMVI\server\etc\cred folder into the C:\Program Files\IBM\Virtual Storage Console\smvi\server\etc\cred folder. After you finish copying the data, you must restart the VSC for VMware vSphere service.

Installing VSC for VMware vSphere using silent mode

You can install VSC for VMware vSphere using silent mode instead of the installation wizard. When you use silent mode, you can automatically install all the capabilities at once.

Before you begin

You must be logged on with administrator privileges to the machine where you are installing VSC for VMware vSphere. In addition, you should make sure you are using the correct installer package for your host machine. You cannot run the 32-bit installer on a 64-bit host machine or the 64-bit installer on a 32-bit host machine.

Procedure

1. Use the following command format to install VSC for VMware vSphere:

```
installer.exe /s /v"/qn /Li logfile ADDLOCAL=ALL INSTALLDIR= \"installation path\"  
""
```

This command installs all the VSC for VMware vSphere capabilities.

Example

The following is a sample command line for a 64-bit host machine:

```
VSC-4.1-win64.exe /s /v"/qn /Li install.log ADDLOCAL=ALL  
INSTALLDIR="C:\Program Files\IBM\Virtual Storage Console\""
```

2. A web page appears when the installation is complete so that you can register vCenter Server VSC for VMware vSphere with the vCenter Server. You must provide the vCenter Server host name or IP address and the administrative credentials.

What to do next

If you uninstall SMVI prior to installing VSC for VMware vSphere, then, before you migrate the backup metadata, you must save this metadata by copying the contents installed in the C:\Program Files\IBM\SMVI\server\repository folder into the C:\Program Files\IBM\Virtual Storage Console\smvi\server\repository folder.

After you finish copying the backup metadata, you must copy the contents of the credential file from the C:\Program Files\IBM\SMVI\server\etc\cred folder into the C:\Program Files\IBM\Virtual Storage Console\smvi\server\etc\cred folder.

After you finish copying the data, you must restart the VSC for VMware vSphere service.

Registering VSC for VMware vSphere with vCenter Server

After installing the VSC for VMware vSphere software, you must register it with the vCenter Server. By default, the registration web page opens when the VSC for VMware vSphere installation is complete.

Procedure

1. If the registration web page does not open automatically, point a web browser to the following URL: `https://localhost:8143/Register.html`. If you use a different computer from the one where you installed VSC for VMware vSphere, replace `localhost` with the hostname or IP address of the computer where you installed VSC for VMware vSphere.

If a security certificate warning appears, choose the option to ignore it or to continue to the web site.

2. In the plug-in service information section, select the IP address the vCenter Server uses to access VSC for VMware vSphere. This IP address must be accessible from the vCenter Server. If you installed VSC for VMware vSphere on the vCenter Server computer, this might be the same address as you use to access the vCenter Server.

Note: IPv6 addresses are not currently supported.

3. Type the host name or IP address of the vCenter Server and the administrative credentials for the vCenter Server.
4. Click **Register** to complete the registration.

Note: A registration failed error message displays if you type the incorrect user credentials for the vCenter Server.

Note: You should close the registration page after you complete the registration process because the web page is not automatically refreshed.

What to do next

If you registered VSC for VMware vSphere with an incorrect vCenter Server, you can register VSC for VMware vSphere with the new vCenter Server.

VSC for VMware vSphere port requirements

By default, VSC for VMware vSphere uses designated ports to enable communication between its components, which include storage systems and the VMware vCenter server. If you have firewalls enabled, you must ensure that the firewalls are set to allow exceptions.

For firewalls other than Windows, you must manually grant access to specific ports that VSC for VMware vSphere uses. If you do not grant access to these ports, an error message such as Unable to communicate with the server appears.

VSC for VMware vSphere uses the following default ports:

Default port number	Description
443	The VMware vCenter Server and the storage systems listen for secure communications using secure HTTP (SSL) on this port.
80	The VMware vCenter server and the storage systems listen for standard, unencrypted communication via standard HTTP on this port.
8143	VSC for VMware vSphere listens for secure communication on this port.
8043	The Backup and Recovery capability listens for secure communication on this port. The Backup and Recovery Restore Agent and CLI also use this port.

Regenerating an SSL certificate for VSC for VMware vSphere

The SSL certificate is generated when you install VSC for VMware vSphere. The distinguished name (DN) generated for the SSL certificate might not be a common name (CN) ("IBM") that the client machines recognize. By changing the keystore and private key passwords, you can regenerate the certificate and create a site-specific certificate.

Procedure

1. Before you generate a new certificate, you should stop the Virtual Storage Console for VMware vSphere (vsc) service. There are several ways to do this. One way to stop the service is to use the Windows Services control panel.
2. Connect to the Windows console session or the Windows PowerShell console.
3. Go to the VSC for VMware vSphere installation directory and enter the following command: `bin\nvpf ssl setup -cn <HOST>` For `<HOST>`, enter the host name of the system running VSC for VMware vSphere or a fully-qualified domain name of the system running VSC for VMware vSphere.

Example

The following example executes the command from the installation directory and uses host called ESXiTester:

```
C:\Program Files\ IBM\Virtual Storage Console>bin\vsc ssl setup -cn ESXiTester
```

4. At the prompt, enter the default keystore password: `changeit` You will also be prompted to enter a password for the private key (this can be any string you choose).

Note: If you do not have Java JRE version 1.6.0_21 installed in Program Files, your command will look different.

The following files are generated:

- keystore file (default: `etc\nvpf.keystore`)

This is the JKS keystore file.

- keystore properties (default: `etc\keystore.properties`)

This file contains the keystore file path and the keystore and key passwords. The administrator should secure this file and specify `http.ssl.keystore.properties` in `etc\nvpf.override` if the keystore properties file needs to be moved.

5. If you are using the Provisioning and Cloning capability or the Optimization and Migration capability, perform the following two steps:

- a. Change to the VSC for VMware vSphere installation directory and enter the following command: `keytool -export -alias nvpf -keystore nvpf.keystore -file nvpf.cer`

Example

The following example executes the command from the etc directory in the installation directory: **C:\Program Files\IBM\Virtual Storage Console\etc>keytool -export -alias nvpf -keystore nvpf.keystore -file nvpf.cer**

The command creates a new file called `nvpf.cer`.

- b. Import the certificate to the local Java keystore by entering the command:
`c:\Program Files\IBM\Virtual Storage Console\etc>keytool -import -alias nvpf -file nvpf.cer -keystore "c:\Program Files\Java\jdk1.6.0_21\jre\lib\security\cacerts"`

What to do next

- You must secure the `etc\keystore.properties` file and then restart the **vsc** service.

There are several ways to do this, including the following:

- If the installation directory is on a network share directory, move the file to local storage.
- Move the file to storage accessible only to the SYSTEM user, which keeps unauthorized users from being able to view or modify the file.
- You can review and accept the SSL certificate after the vSphere Client receives the certificate when you click the IBM N series icon in the vSphere Client.

You can then import the SSL certificate into the Trusted Root Certification Authorities store to prevent the SSL security warnings from appearing every time you launch the vSphere client. For details, see the documentation for your Windows operating system.

How to configure role-based access control for VSC for VMware vSphere

VSC for VMware vSphere works with the role-based access control (RBAC) feature of Data ONTAP.

As a result, you set up the RBAC information by performing the following tasks from within Data ONTAP:

- Create the roles.
- Assign to these roles the privileges that you want to be available in VSC for VMware vSphere.
- Create user name/password logins in Data ONTAP associated with those roles.

After the logins have been set up in Data ONTAP, you can use the Monitoring and Host Configuration capability to configure the storage system to use the logins. The Monitoring and Host Configuration capability manages credentials for itself, the Provisioning and Cloning capability, and the Optimization and Migration capability, so each time you log into a storage system from one of these three capabilities, you are presented with the same set of VSC for VMware vSphere functions for that user name and password pair that you set up in Data ONTAP.

The Monitoring and Host Configuration capability performs an initial privilege validation for RBAC when the storage system is connected to an administrative

Vserver. The Monitoring and Host Configuration capability also performs this upfront validation for systems using pFiler units. It does not perform this validation if the storage system is directly connected to a Vserver or a vFiler unit.

Note: Both the Provisioning and Cloning capability and the Optimization and Migration capability support additional restrictions provided by the vCenter server. The Provisioning and Cloning capability lets you use the vCenter Server Provisioning and Cloning privileges, and the Optimization and Migration capability lets you use the vCenter Server Optimization and Migration privileges.

The Monitoring and Host Configuration capability performs an initial privilege validation for RBAC when the storage system is connected to an administrative Vserver. The Monitoring and Host Configuration capability also performs this upfront validation for systems using pFiler units. It does not perform this validation if the storage system is directly connected to a Vserver or a vFiler unit.

The Backup and Recovery capability manages its own credentials. While you can use the same logins that you created in Data ONTAP, you must add them separately to the Backup and Recovery capability. That capability does not recognize logins that you add using the Monitoring and Host Configuration capability. In addition, any logins that you set up for the Backup and Recovery capability apply only to that capability.

Note: You can use the root login for all the capabilities; however, it a good practice to use the RBAC feature provided by Data ONTAP to create one or more custom accounts with limited functions.

Related concepts:

“Authentication methods in the Backup and Recovery capability” on page 98

Preparation required before upgrading VSC for VMware vSphere

Using shared storage system credentials for the Monitoring and Host Configuration, Provisioning and Cloning, and Optimization and Migration capabilities changes how storage systems are added to VSC for VMware vSphere and which credentials are used. This upgrade process also overwrites the preferences file with a new file. As a result, you should check information about your storage systems, the credentials being used, and the preferences being used before you upgrade to a new version of VSC for VMware vSphere.

It is a good practice to make a note of your storage resources before you upgrade to 4.0 or later of VSC for VMware vSphere. After the upgrade, the Monitoring and Host Configuration capability exports to the Provisioning and Cloning and Optimization and Migration capabilities all the storage systems with valid credentials that were either automatically discovered by or manually added to the Monitoring and Host Configuration capability. Each storage system now uses the same credentials in all three capabilities.

The Monitoring and Host Configuration capability manages the credentials for all three capabilities so you enter the credentials at the Overview panel for the Monitoring and Host Configuration capability. That capability then exports the information to the Provisioning and Cloning and Optimization and Migration capabilities.

After an upgrade, you might discover that some of the storage systems that were previously configured in the Provisioning and Cloning capability are no longer

there. This situation occurs when storage systems were added to the Provisioning and Cloning capability, but not discovered or added to the Monitoring and Host Configuration capability. In these cases, you must either get the Monitoring and Host Configuration capability to discover the storage systems or manually add them to the Monitoring and Host Configuration capability. The Monitoring and Host Configuration capability then exports the storage systems to the Provisioning and Cloning and Optimization and Migration capabilities.

Note: If the storage system does not have storage mapped to an ESX/ESXi host that a vCenter Server is managing, the Monitoring and Host Configuration capability does not automatically discover it.

Another point to keep in mind is that, with shared credentials, all the storage systems use the same credentials. You cannot have one set of credentials for Provisioning and Cloning capability and a different set for Optimization and Migration capability. Both the Provisioning and Cloning capability and the Optimization and Migration capability support additional restrictions provided by the vCenter server. The Provisioning and Cloning capability lets you use the vCenter Server Provisioning and Cloning privileges, and the Optimization and Migration capability lets you use the vCenter Server Optimization and Migration privileges.

In addition, if you made changes to the preferences file for any of the capabilities in the 2.1.x version of VSC for VMware vSphere, you should record those changes so that you can enter them again after you install a 4.0 or later version because the upgrade process does not record them.

Upgrading VSC for VMware vSphere

VSC for VMware vSphere supports upgrades from version 2.0 or later of VSC for VMware vSphere. The VSC for VMware vSphere installer checks the version numbers of each of the currently installed capabilities to determine whether you are upgrading to a newer version.

Before you begin

If you are upgrading from version 2.x of VSC for VMware vSphere to version 4.0 or later, the VSC for VMware vSphere installer automatically upgrades the Monitoring and Host Configuration, Optimization and Migration, and Provisioning and Cloning capabilities to the newer versions. If you also have the Backup and Recovery capability installed, the VSC for VMware vSphere installer upgrades it as well. If you do not have the Backup and Recovery capability installed, the VSC for VMware vSphere installer gives you the option of installing it.

The VSC for VMware vSphere installer does not support upgrades from the following:

- A version of VSC for VMware vSphere prior to 2.0
- A standalone version of Rapid Cloning Utility (RCU), which is now called the Provisioning and Cloning capability
- A standalone version of SnapManager for Virtual Infrastructure (SMVI), which is now called Backup and Recovery.

If you have any of that software installed, you must uninstall it before you can install the current version of VSC for VMware vSphere. If the VSC for VMware

vSphere installer finds one of those versions of VSC for VMware vSphere, RCU, or SMVI on the server, it prompts you to uninstall the software, and then aborts.

You must be logged on with administrator privileges to the machine where you installing VSC for VMware vSphere.

Procedure

1. Obtain the product software by downloading the software as follows:
 - a. Go to the IBM NAS support Web site.
 - b. Sign in with your IBM ID and password. If you do not have an IBM ID or password, or if you are adding new N series machines and serial numbers to an existing registration, click the **Register** link, follow the online instructions, and then sign in.
 - c. Select the N series software you want to download, and then select the **Download** view.
 - d. Use the **Software Packages** link on the web page presented and follow the online instructions to download the software to a working directory on the Windows server.
2. Double-click the installer icon, and click **Run** to start the installation wizard.
3. Click **Yes** on the confirmation prompt.
4. In the installation wizard, select the capabilities that you want to upgrade and click **Next** to start the installation.

The installation might take several minutes.
5. Click **Finish** to complete the installation. A web page appears when the installation is complete. You must register VSC for VMware vSphere with the vCenter Server. You must provide the vCenter Server host name or IP address and the administrative credentials.

Adding or removing the Backup and Recovery capability

After you have installed the VSC for VMware vSphere software, you can use the installation wizard to add or remove the Backup and Recovery capability.

About this task

Unlike the other capabilities, the Backup and Recovery capability is not automatically installed when you run the installation wizard. You must select the Backup and Recovery checkbox to install it. As a result, you can use this checkbox to uninstall only the Backup and Recovery capability.

Note: Backup and Recovery is the only capability you can uninstall or install separately. The other three capabilities can only be installed or uninstalled as a group.

Procedure

1. On the Windows server where you installed the VSC for VMware vSphere software, select **Control Panel > Add/Remove Programs** (Windows Server 2003) or **Control Panel > Programs and Features** (Windows Server 2008).
2. Select Virtual Storage Console for VMware vSphere and click **Change** to start the installation wizard.
3. In the installation wizard, select the **Modify** option and click **Next**.

4. Select the checkbox for the Backup and Recovery capability if you want to add it or clear the checkbox if you want to remove it. Now click **Next**.
5. If you are installing the capability, click **Install** to start the installation. The installation might take several minutes.
6. Click **Finish** to complete the installation or removal of the Backup and Recovery capability.

What to do next

You must close the vSphere Client and restart it to either display any newly-installed capability or remove it from the GUI.

Related concepts:

“Backing up and restoring data” on page 97

Uninstalling VSC for VMware vSphere using Add/Remove Programs

You can uninstall the VSC for VMware vSphere software from your system using the Windows Add or Remove Programs list.

About this task

The uninstall program removes the entire VSC for VMware vSphere application. You cannot specify which capabilities you want to uninstall.

Procedure

1. On the Windows server where you installed the VSC for VMware vSphere software, select **Control Panel > Add/Remove Programs** (Windows Server 2003) or **Control Panel > Programs and Features** (Windows Server 2008).
2. Select Virtual Storage Console for VMware vSphere and click **Remove** to immediately remove the program or click **Change** to start the installation wizard.
3. If you select **Change**, then click **Yes** to confirm that you want to remove the program.
4. In the installation wizard, select the **Remove** option and click **Next**.
5. Click **Remove** to uninstall the VSC for VMware vSphere software. After the process completes, a confirmation prompt is displayed.

Note: At the confirmation prompt, click **Yes** to remove all the metadata files from the installation directory or click **No** so that you can manually delete the files in the directory.

Uninstalling VSC for VMware vSphere using silent mode

You can uninstall VSC for VMware vSphere using silent mode instead of the Windows Add/Remove Program. With the command line, you can automatically uninstall all the capabilities at once.

Before you begin

You must be logged on with administrator privileges to the machine where you are uninstalling VSC for VMware vSphere.

Procedure

Use the following command format to uninstall VSC for VMware vSphere:
installer.exe /s /v"/qn /Li logfile REMOVE=ALL INSTALLDIR= \"installation path\" ""
This command removes all the VSC for VMware vSphere capabilities.

Example

The following is an example of the command line you might use if you were uninstalling the 4.0 version of VSC for VMware vSphere from a 64-bit host machine:

```
VSC-4.0-win64.exe /s /v"/qn /Li uninstall.log REMOVE=ALL
```

Monitoring and configuring hosts and storage controllers

The Monitoring and Host Configuration capability enables you to work with storage controllers and ESX and ESXi hosts. It also manages the task of adding storage controllers and setting the credentials for the Provisioning and Cloning and Optimization and Migration capabilities.

The tasks that you can perform with the Monitoring and Host Configuration capability include the following:

- Add, manage, and remove storage controllers for the Monitoring and Host Configuration, Provisioning and Cloning and Optimization and Migration capabilities.

Note: The Monitoring and Host Configuration capability automatically launches either Element Manager or the FilerView GUI to enable you to manage storage controllers.

- Set credentials to access storage controllers for the Monitoring and Host Configuration, Provisioning and Cloning and Optimization and Migration capabilities.

- View status for the following:

- Storage controllers

You can quickly see the status of storage controllers from the following perspectives:

- SAN (FC, FCoE, and iSCSI)
- NAS (NFS)

Note: The Monitoring and Host Configuration capability displays the status of storage controllers running Data ONTAP operating in Cluster-Mode and 7-Mode as well as other versions of Data ONTAP.

- SAN and NAS datastore capacity utilization
- VMware vStorage APIs for Array Integration (VAAI) support in the storage controller
- ESX and ESXi hosts, including the operating system version and overall status
- Configure the correct values for the following:
 - Storage adapter timeouts
 - Multipathing settings
 - NFS settings
- Determine whether the paths to NFS nodes provide direct and indirect data access to NFS datastores
- Collect diagnostic information from the ESX and ESXi hosts, storage controllers, and Fibre Channel switches
- Access tools to set guest operating system timeouts and to identify and correct misaligned disk partitions

Note: The tools for aligning disk partitions that are provided by the Monitoring and Host Configuration capability require that you power down the VM. You can use the online alignment tool provided by the Optimization and Migration capability to align disk partitions without having to power down the VMs.

- Install and enable the NFS Plug-in for VMware VAAI software library

You perform these tasks from the Monitoring and Host Configuration capability GUI. To access the GUI, you click the IBM N series icon in the vCenter Server and then click Monitoring and Host Configuration in the navigation pane.

VSC for VMware vSphere configuration

You can configure and manage your ESX and ESXi hosts and virtual machines (VMs) by first specifying the physical storage systems on which the active images of the datastores and VMs that are managed by the vCenter Server reside.

Storage system discovery and credentials overview

VSC for VMware vSphere uses a single mechanism to discover storage systems for all capabilities except the Backup and Recovery capability. Each capability requires certain Data ONTAP permissions to perform its operations.

Note: The Monitoring and Host Configuration capability manages the storage system credentials for the Optimization and Migration capability and the Provisioning and Cloning capability as well as its own storage system credentials.

Before VSC for VMware vSphere can display and manage storage resources, it must discover the storage systems that provide the storage. As part of the discovery process, you must supply storage system credentials. When the Monitoring and Host Configuration capability adds a storage system, it displays a pop-up box that lists the credentials (or privileges) associated with the username and password pair you entered when you logged in. You can either set up default credentials that Monitoring and Host Configuration capability will use during its discovery, or manually enter credentials when the storage system is discovered.

Note: If you have vFiler units on storage systems running Data ONTAP 8.x software, you must set the `httpd.admin.enable` for the vFiler unit in order to enable discovery.

Discovery happens in one of the following ways. In each case, you must supply credentials for any newly discovered storage system.

- When the VSC for VMware vSphere service starts, the Monitoring and Host Configuration capability begins its automatic background discovery process.
- You click **Update** on the Monitoring and Host Configuration capability Overview panel.

This runs another automatic discovery.

- You open the Add Storage System dialog box in the Setup panel of Backup and Recovery and enter the address of the storage controller and its credentials.

This discovery is manual, not automatic.

Note: IPv6 addresses are not currently supported.

Storage system credentials are assigned based on the user name/password pairs. This can be the root account or a custom account that uses role-based access control (RBAC). Each role consists of a group of privileges. You cannot change the role associated with a user name/password pair at the Modify Credentials dialog box.

The Backup and Recovery capability discovers its own set of storage system and maintains its own list of storage system credentials. The automatic discovery process of the Monitoring and Host Configuration capability does not affect the list of credentials maintained by the Backup and Recovery capability. You can specify the same or different credentials for a given storage system used by both capabilities.

You cannot specify credentials for the Optimization and Migration capability and the Provisioning and Cloning capability. The Monitoring and Host Configuration manages these credentials.

All of the capabilities require specific permissions to perform certain RBAC operations. You can limit what users can do in these capabilities based on the credentials associated with their vSphere Client account. All users of Backup and Recovery share the same set of storage system credentials and can all perform the same operations.

Related tasks:

“Discovering and adding storage resources” on page 24

“Enabling discovery and management of vFiler units” on page 22

“Enabling discovery and management of vFiler units on private networks” on page 23

Default controller credentials simplify administration of capabilities

You can set up default storage controller credentials in the Monitoring and Host Configuration capability of VSC for VMware vSphere. You do not have to manually specify credentials for any storage controller for which the default credentials are valid. These credentials apply to Monitoring and Host Configuration capability, Optimization and Migration capability, and Provisioning and Cloning capability.

Note: The Backup and Recovery capability discovers its own set of storage controllers and maintains its own list of storage controller credentials.

When the Monitoring and Host Configuration capability discovers a new storage controller, it attempts to log in using the default credentials. If the login fails, the controller status is set to Authentication Failure, and you must enter credentials manually from the Overview panel by right-clicking the controller name and choosing **Modify Credentials** from the pop-up menu.

You can set the default credentials by clicking **Set Default Controller Credentials** on the Discovery Status panel of the Monitoring and Host Configuration capability.

Any time you change the default credentials and run **Update**, the Monitoring and Host Configuration capability uses the new credentials and attempts to log in to any controller that has a status of either Authentication Failure or SSL is not configured.

Specifying credentials with Monitoring and Host Configuration

You can use the Monitoring and Host Configuration capability to set up credentials for a storage controller. These credentials then apply to the Monitoring and Host Configuration capability, the Optimization and Migration capability, and the Provisioning and Cloning capability.

About this task

You either set up default credentials for each storage controller, or you can manually assign the credentials.

Procedure

1. Select the Overview panel of the Monitoring and Host Configuration capability.
2. Right-click a storage controller that needs credentials.
3. From the pop-up dialog menu that appears, select **Modify Credentials**.
4. Fill in the following information:
 - Management IP address
VSC for VMware vSphere uses the management IP address to communicate with the controller. VSC for VMware vSphere lists the available addresses.
 - Management port number
The default management port number is 443 if the SSL box is checked and 80 if it is not checked. These are the Data ONTAP defaults. If you toggle the SSL checkbox, the port number switches between 443 and 80. You can specify a different port number. If you do that, then toggling the SSL check box only changes the SSL state in the dialog box.
 - Whether SSL is enabled
 - User name/password
Storage controller credentials are assigned based on the user name/password pair. This can be the root account or a custom account that uses role-based access control (RBAC). You cannot change the roles associated with that user name/password pair at the Modify Credentials dialog box.
 - Whether the controller is skipped
If you chose to **not** provide credentials for this controller, you must select the Skipped check box. Clear the check box and enter the credentials for the controller.

If a controller is skipped, the Monitoring and Host Configuration capability does not export it to the Provisioning and Cloning and Optimization and Migration capabilities. If it exists for these capabilities, it is deleted from their controller lists and unavailable for their work flows.

However, if you uncheck the Skipped check box, the Monitoring and Host Configuration capability adds the controller back to the export list and the Optimization and Migration capability and the Provisioning and Cloning capability can add it to the work flows again.
5. When you have filled out the information in the dialog box, click **OK**. VSC for VMware vSphere displays the list of allowed and disallowed roles.
6. If the allowed roles support your needs, then click **OK**. If the roles are not sufficient, click **Cancel**. Doing this returns you to the **Modify Credentials**. You can either enter a different user name/password pair that will provide different credentials or talk with your system administration about them.

Once a storage controller has credentials, you can view the privileges by right-clicking the controller name on the Overview panel.

Enabling discovery and management of vFiler units

If you are using Data ONTAP 8, you must set the `httpd.admin.enable` option for vFiler units in order to enable discovery and management with the VSC for VMware vSphere.

About this task

This task is not required for vFiler units created with Data ONTAP 7.x.

Procedure

1. From the storage system, enter the following command to switch to a particular vFiler context: `vfiler context vfiler_name`
2. Enter the following command in the vFiler context to set the required option that enables discovery in VSC for VMware vSphere: `options httpd.admin.enable on`
3. Repeat for each vFiler unit you want to manage using VSC for VMware vSphere.

Enabling discovery and management of vFiler units on private networks

If vFiler units are isolated in private networks to which the VSC for VMware vSphere has no network connectivity, you must manually add the parent `vfiler0` to the Monitoring and Host Configuration capability.

Before you begin

The VSC for VMware vSphere server must have network connectivity to the parent `vfiler0`.

Procedure

1. On the Monitoring and Host Configuration capability Overview panel, right-click within the Storage Controllers section and select **Add Controller**.
2. Enter the management IP address and credentials for the parent `vfiler0` and then click **OK**.

Results

Any vFiler units belonging to the parent `vfiler0` that provide storage to ESX hosts are discovered by the Monitoring and Host Configuration capability.

Tunneled vFiler units and Vservers discovered automatically

The Monitoring and Host Configuration capability automatically supports vFiler and Vserver tunneling for the storage systems it manages. You do not need to manually add these vFiler units and Vservers to the Monitoring and Host Configuration capability.

When you enter information for a cluster administrative LIF or `vfiler0`, the Monitoring and Host Configuration capability discovers all the subordinate vFiler units and Vservers..

Differences between direct connections to Vservers and to cluster-management LIFs

Virtual Storage Console for VMware vSphere supports connecting a storage controller directly to either a Vserver or a cluster-management LIFs. When the storage connects directly to a Vserver, not all of the Virtual Storage Console for VMware vSphere features are supported. To use all the features, you must connect the storage to a cluster-management LIF.

Virtual Storage Console for VMware vSphere does not provide the following features when the storage controller connects directly to a Vserver:

- Upfront validation of Role-Based Access Control (RBAC)
While RBAC is fully supported, Virtual Storage Console for VMware vSphere does not perform the initial privilege validation on storage that is directly connected to a Vserver.
- NFS path checking
When you are running Data ONTAP operating in Cluster-Mode and using a cluster-management LIF, the Monitoring and Host Configuration capability can query the storage controller to determine whether the storage controller is using a direct or indirect path. The Monitoring and Host Configuration capability then reports this information, and supplies information you can use to set up a direct path. Better performance is normally seen when direct paths are used. If a storage controller connects directly to a Vserver, the Monitoring and Host Configuration capability cannot query the storage controller to determine the path.
- Reports on space that is shared by volumes using data deduplication
Virtual Storage Console for VMware vSphere is not able to check the space shared by volumes that have data deduplication enabled when the storage controller is directly attached to a Vserver.
- EMS logging
Virtual Storage Console for VMware vSphere cannot perform EMS logging when the storage controller is directly attached to a Vserver.
- Storage-side log collections for the nSANity Diagnostic and Configuration Data Collector program
The nSANity program collects data that can be used to resolve problems. You should only use the nSANity program if you are directed to do so by technical support. When the storage connects directly to a Vserver, the nSANity program cannot collect information on that storage.

These features are supported when the storage controller connects to a cluster-management LIF.

Discovering and adding storage resources

When you first run VSC for VMware vSphere in a vSphere Client, the Monitoring and Host Configuration capability discovers ESX and ESXi hosts, their LUNs and NFS exports, and the IBM N series storage systems that own those LUNs and exports. You must provide the storage system credentials.

About this task

You can discover new resources and get the latest capacity and configuration information at any time by clicking **Update** on the Overview panel of the Monitoring and Host Configuration capability in the vSphere Client.

The discovery process collects information from the ESX and ESXi hosts managed by the vCenter Server. Make sure all ESX and ESXi hosts are shown as powered on and connected.

This discovery process is for the Monitoring and Host Configuration, Optimization and Migration, and Provisioning and Cloning capabilities only. You must add storage controllers to the Backup and Recovery capability manually.

Procedure

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the Inventory panel, and then select the IBM N series tab.
3. If the discovery process does not start automatically, or if you want to discover new resources and update information, click **Update** on the Overview panel of the Monitoring and Host Configuration capability.
4. Right-click any discovered storage controllers with the status Authentication Failure and select **Modify Credentials**.
5. Fill in the information in the **Modify Credentials** dialog box.

What to do next

After discovery is complete, use the Monitoring and Host Configuration capability to configure ESX or ESXi host settings for any hosts displaying an Alert icon in the Adapter Settings, MPIO Settings, or NFS Settings columns.

Related concepts:

"Storage system discovery and credentials overview" on page 20

Related tasks:

"Getting information about an ESX and ESXi host with an Alert status" on page 168

Correcting controller names displayed as "unknown"

If the Monitoring and Host Configuration capability displays a controller name as "unknown" on the Overview panel, you can modify the credentials and add the management IP address of the controller in the Modify storage system -unknown- pop-up box.

About this task

This issue can occur if an NFS datastore is mounted over a private network.

If you are running Data ONTAP operating in Cluster-Mode and working with NFS datastores that are mounted using an NFS data LIF, this issue can occur with either a private network or a public network.

Procedure

1. Right-click the controller and select **Modify Credentials**.
2. Enter the management IP address of the storage controller and the storage controller credentials in the Modify storage system -unknown- pop-up box. VSC for VMware vSphere must have network connectivity to the management port you specify.

Removing controllers from the Monitoring and Host Configuration capability

You can remove a skipped or unmanaged storage controller that is not attached to a host. When you remove a storage controller, it no longer appears in the Monitoring and Host Configuration, Provisioning and Cloning, or Optimization and Migration display.

About this task

If a storage controller has storage mapped to an ESX or ESXi host managed by the Monitoring and Host Configuration capability and you attempt to remove that storage controller, the Monitoring and Host Configuration capability displays an error message and does not remove the storage controller. You can only remove storage controllers that are not attached to hosts.

Procedure

1. Open the vSphere Client and log in to your vCenter Server.
2. Select a Datacenter in the Inventory panel, and then select the IBM N series tab.
3. Select Monitoring and Host Configuration and then select the Overview panel.
4. Right-click the storage controller and select **Remove Controller**.

If the storage controller is ...	The Monitoring and Host Configuration capability
Not attached to host	Removes the storage controller.
Attached to a host	Displays an error message and does not change the controller

Administering the Monitoring and Host Configuration capability

The Monitoring and Host Configuration capability lets you work with hosts and controllers and configure values for them.

The sections that follow apply to using the Monitoring and Host Configuration capability.

Inventory panel selection limits what VSC for VMware vSphere displays

The Monitoring and Host Configuration capability displays only the resources associated with what is selected in the vSphere Client Inventory panel. If you click the **Update** button, the Monitoring and Host Capability updates only those selected resources.

You can access the Monitoring and Host Configuration panels in two ways. If you select the **IBM N series** icon in the **Solutions and Applications** section of the vSphere Client **Home** page, the Inventory panel is not displayed, and the Monitoring and Host Configuration panels display all discovered resources.

If instead you access the Monitoring and Host Configuration panels using the IBM N series tab in the **Inventory** section of the vSphere Client, the Monitoring and Host Configuration panels display only those resources associated with the selection in the Inventory panel. To display all resources, you must select the Datacenter object in the Inventory panel.

If you cannot find an expected resource in one of the Monitoring and Host Configuration panels, the first thing to check is the selection in the Inventory panel.

The effect of the **Update** button is also limited by the selection in the Inventory panel. Because actions are limited to what is selected in the Inventory panel (a

subset of the entire configuration), you get faster results. To update the entire configuration, you need to select the Datacenter object in the Inventory panel before clicking **Update**.

Configuring ESX server multipathing and timeout settings

The Monitoring and Host Configuration capability checks and sets the ESX or ESXi host multipathing and HBA timeout settings that ensure proper behavior with IBM N series storage systems.

Before you begin

This process might take a long time, depending on your configuration and system load. The task progress is displayed in the **Recent Tasks** panel. As tasks complete, the host status Alert icons are replaced by Normal or Pending Reboot icons.

Procedure

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the Inventory panel, and then select the IBM N series tab.
3. Select the **Overview** panel.
4. Select one or more ESX hosts that have an Alert icon in the **Adapter Settings**, **MPIO Settings**, or **NFS Settings** columns. Use Ctrl-click or Shift-click to select multiple hosts.
5. Right-click the selected hosts and select **Set Recommended Values**.
6. Select the types of settings you want to update and then click **OK**.

ESX host settings set by Monitoring and Host Configuration capability

Monitoring and Host Configuration capability sets ESX or ESXi host timeouts and other settings to ensure best performance and successful failover.

Monitoring and Host Configuration capability sets the following values on an ESX or ESXi host.

NFS Settings

Net.TcpipHeapSize

If you are using vSphere 5.0 or later, set to 32.

For all other NFS configurations, set to 30.

Net.TcpipHeapMax

If you are using vSphere 5.0 or later, set to 128.

For all other NFS configurations, set to 120.

NFS.MaxVolumes

If you are using vSphere 5.0 or later, set to 256.

For all other NFS configurations, set to 64.

NFS.HeartbeatMaxFailures

Set to 10 for all NFS configurations.

NFS.HeartbeatFrequency

Set to 12 for all NFS configurations.

NFS.HeartbeatTimeout

Set to 5 for all NFS configurations.

FC/FCoE Settings

Path selection policy

Set to RR (round robin) for ESX 4.0 or 4.1 and ESXi 5.0, FC paths with ALUA enabled. Set to FIXED for all other configurations.

Disk.QFullSampleSize

Set to 32 for all configurations. This setting is available with ESXi 5.0 and ESX 4.x.

Note: vSphere 5.1 handles Task Set Full (QFull) conditions differently from vSphere 4.x and 5.0.

Disk.QFullThreshold

Set to 8 for all configurations. This setting is available with ESXi 5.0 and ESX 4.x.

Note: vSphere 5.1 handles Task Set Full (QFull) conditions differently from vSphere 4.x and 5.0.

Emulex FC HBA timeouts

For ESX 4.0 or 4.1 or ESXi 5.0 or later, use the default value.

QLogic FC HBA timeouts

For ESX 4.0 or 4.1 or ESXi 5.0 or later, use the default value.

iSCSI Settings

Path selection policy

Set to RR (round robin) for all iSCSI paths.

Disk.QFullSampleSize

Set to 32 for all configurations. This setting is available with ESX 4.x and ESXi 5.0.

Note: vSphere 5.1 handles Task Set Full (QFull) conditions differently from vSphere 4.x and 5.0.

Disk.QFullThreshold

Set to 8 for all configurations. This setting is available with ESX 4.x and ESXi 5.0.

Note: vSphere 5.1 handles Task Set Full (QFull) conditions differently from vSphere 4.x and 5.0.

QLogic iSCSI HBA IP_ARP_Redirect

Set to ON for all configurations.

QLogic iSCSI HBA timeouts

ql4xportdownretrycount (qla4022 driver), ka_timeout (qla4xxx driver), and KeepAliveTO timeout settings are set to 14 for iSCSI SAN booted ESX hosts, and set to 60 for non-SAN-boot configurations.

UNMAP setting turned off in ESX 5.x

On hosts running ESX 5.0, the Monitoring and Host Configuration capability automatically turns off the UNMAP (`VMFS3.EnableBlockDelete`) parameter by setting it to 0.

On ESX 5.1 hosts, 0 is the default value. If you change this value to 1 on either ESX 5.0 or 5.1, the Monitoring and Host Configuration capability automatically resets it to 0. For this value to take effect, you must apply the HBA/CNA Adapter Settings on the host.

To avoid any potential performance impact due to UNMAP operations, VMware disabled this feature beginning in ESXi 5.0 Patch 2. VSC 4.1 for VMware vSphere ensures that this feature is disabled with all versions of ESXi 5.x.

Direct path access and NFS datastores

When you are running Data ONTAP operating in Cluster-Mode, it is possible for a client to access a data LIF with an indirect data path to a FlexVol. Indirect data paths can affect I/O performance and should be corrected. The Monitoring and Host Configuration capability provides tools to scan for direct and indirect NFS paths and provide you with the information you need to manually correct the paths.

This situation can occur when a data LIF is bound to a different physical node than the one that owns the exported FlexVol. The NFS virtual client does not have the path selection intelligence that is native to physical clients. In order to have a direct data path, the client must access a data LIF that is local to the node that owns the exported FlexVol.

The Monitoring and Host Configuration capability monitors which LIFs NFS is using to access the volume. You can see whether a LIF uses a direct data path or an indirect data path by going to the Storage Details - NAS panel fields and commands window and viewing the Data Path Access column. This column displays the path setting as Direct (green check), Indirect (red exclamation point (!)), N/A, or (unknown).

If the path setting is indirect, you can right-click that row and select the **View Direct Data Path Choices** option. This option displays the **Direct Data Path Choices** pop-up, which contains a list of ports using direct paths to access data.

Note: VSC for VMware vSphere does not check these ports to ensure that they are connected to the network. You must do that manually.

Anytime the data path access changes, either to direct from indirect or to indirect from direct, the Monitoring and Host Configuration capability writes the path information to a log file.

If a direct Vserver connection is made, the Monitoring and Host Configuration capability cannot query the storage controller to determine the path.

An N/A (not applicable) entry indicates a path to a storage controller running Data ONTAP operating in 7-Mode, so there is not an issue about whether the path is direct.

An unknown path occurs if the discovery data is incomplete.

Changing NFS data paths to direct access

If you have a cluster node that is accessing a data LIF with an indirect data path, you can change the path to one that is direct.

Before you begin

Only a storage administrator should change the path.

About this task

This issue only occurs if you are running Data ONTAP operating in Cluster-Mode and you have an NFS datastore using a remote data LIF that is bound to a different physical node than the one that owns the exported FlexVol.

Procedure

1. In the Storage Details - NAS panel fields and commands window, right-click a row that has an indirect data path (shown as Indirect) and select the **View Direct Data Path Choices** option. This option displays the **Direct Data Path Choices** pop-up, which contains a list of ports to data paths providing direct access. You cannot use this window to change the path, but you can use it to get information about the available ports.
2. Manually check to make sure the port you want to use is connected to the network. The Monitoring and Host Configuration capability displays the ports without checking their network connectivity. If you try to use a port that is not connected to the network, your datastore will go offline.
3. Once you have confirmed that the path you want to use is connected to the network, collect the information displayed in the **Direct Data Path Choices** pop-up and give it to a storage administrator. Only the storage administrator should change the path. The **Direct Data Path Choices** pop-up contains all the other information a storage administrator needs to move the LIF.

To create a data path with direct access, you must have the correct credentials.

Note: If multiple datastores are using that LIF, moving the LIF will cause the other datastores to have data paths with indirect data access.

4. Use either the storage controller console or a tool such as System Manager to change the path.

Note: Anytime the path value changes, the Monitoring and Host Configuration capability writes the information to a log file.

NFS Plug-in for VMware VAAI requires additional installation steps

The NFS Plug-in for VMware VAAI is not shipped with VSC for VMware vSphere; however, you can get it from the N series support website (accessed and navigated as described in Websites) and then use the Monitoring and Host Configuration capability to install it.

The plug-in is supported on systems running ESXi 5.0 or later with vSphere 5.0 and Data ONTAP 8.1 or later operating in Cluster-Mode or Data ONTAP 8.1.1 or later operating in 7-Mode.

After you install the plug-in, you must reboot the host. To remind you of this, the Monitoring and Host Configuration capability changes the status of the host in the Overview panel to **Pending Reboot**.

For the most up-to-date information about the plug-in, see the *Release Notes*.

Using Web-based tools to manage storage

Data ONTAP provides tools that you can use to create LUNs and manage storage systems. These are Web-based tools that you can use to manage common storage system functions from a Web browser. Depending on the tool, Monitoring and Host Configuration capability either launches it or points you to the download location for the tool.

Before you begin

The tool that the Monitoring and Host Configuration capability displays or recommends depends on which version of Data ONTAP you are running on the storage system:

- For storage systems running Data ONTAP 8.1 or later and operating in 7-Mode, the Monitoring and Host Configuration capability supports OnCommand System Manager.

The Monitoring and Host Configuration capability displays a pop-up box that contains a link to the N series support website (accessed and navigated as described in Websites) where you can download OnCommand System Manager. The Monitoring and Host Configuration capability displays this pop-up box even if you already have OnCommand System Manager installed. It does not launch OnCommand System Manager or install OnCommand System Manager.

- For storage systems running Data ONTAP operating in Cluster-Mode, the Monitoring and Host Configuration capability lets you launch Element Manager.
- For storage controllers running Data ONTAP 8.0 and earlier and operating in 7-Mode, the Monitoring and Host Configuration lets you launch FilerView.

Note: FilerView is not supported with vFiler units, storage systems running Data ONTAP operating in Cluster-Mode, or Data ONTAP 8.1 and later.

Your client computer must have network connectivity to the management port of a storage system to be able to run the tool for that storage system.

Procedure

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the Inventory panel, and then select the IBM N series tab.
3. Select the Monitoring and Host Configuration capability **Overview** panel.
4. Right-click the storage system you want to manage and the Monitoring and Host Configuration capability allows you to select either **Open Element Manager** or **Open FilerView**.

The Monitoring and Host Configuration capability automatically displays the appropriate tool for your version of Data ONTAP. If you select **Open FilerView** and the storage system is running Data ONTAP 8.1 or later, the Monitoring and Host Configuration capability displays a pop-up box recommending that you download OnCommand System Manager because FilerView is not available on those releases. The Monitoring and Host Configuration capability does not provide OnCommand System Manager.

Note: If you using a version of Data ONTAP operating in 7-Mode that supports FilerView and you set up your storage to use a Secure Sockets Layer (SSL) protocol, FilerView automatically opens using the HTTPS protocol. Otherwise it opens using the HTTP protocol.

5. If the storage system requires a password, enter the user name and password when prompted.

Results

For versions of Data ONTAP that support them, either Element Manager or the FilerView GUI opens in a new browser window. You can then use the tool to work with LUNs on that storage system. For storage systems using Data ONTAP 8.1 or later, you must download OnCommand System Manager and use that to work with LUNs.

The Monitoring and Host Configuration capability displays configuration details

You can display details of your storage configuration using the Monitoring and Host Configuration capability of the VSC for VMware vSphere.

The Monitoring and Host Configuration capability displays four pages of configuration information.

Note: See the online help explanations of the fields displayed and available commands.

Overview

Displays the status of storage controllers and ESX and ESXi hosts. You can right-click a host to get detailed configuration information.

Storage Details - SAN

Displays information about the VMFS datastores. If you select a datastore, you can see a list of all LUNs mapped to managed ESX and ESXi hosts. When you select a LUN, the LUN details are displayed.

Storage Details - NAS

Displays a list of all NFS exports mounted on managed ESX and ESXi hosts. When you select a datastore, the details are displayed.

Discovery Status

Displays details of the storage resources discovered by the Monitoring and Host Configuration capability.

Related concepts:

"Online help" on page 6

MultiStore vFiler units are displayed differently

The Monitoring and Host Configuration capability displays vFiler units differently than physical storage controllers.

When you have a vFiler unit create with the optional MultiStore feature of Data ONTAP software, the Monitoring and Host Configuration capability displays the following information:

- The hostname displays a "MultiStore" prefix to identify vFiler units.
- The **Supported Protocols** column reports the storage protocols actually in use by ESX and ESXi hosts instead of the protocols licensed for the storage controller.
- The **Alert** icon in the **Status** column means that the vFiler unit does not respond to the Monitoring and Host Configuration capability. The **Normal** icon means that the Monitoring and Host Configuration capability is able to communicate with the vFiler unit.
- No detailed status is returned for vFiler units. The **Status Reason** column displays This controller is a MultiStore vFiler unit. You can connect to the physical controller that owns the vFiler unit to get more status information.

- FilerView is not available for vFiler units. The **Open FilerView** menu item is not shown when you right-click a vFiler unit.
- In the Storage Details screen for vFiler units, no aggregate information is displayed.
Direct vFiler units and direct vServers do not have aggregates to display.

Collecting diagnostic information

You can use the Monitoring and Host Configuration capability to collect diagnostic information about storage controllers, ESX and ESXi hosts, and Fibre Channel switches.

Before you begin

If VSC for VMware vSphere is running on a Windows Server 2008 or Server 2008 R2 system, the VSC for VMware vSphere service must run on an Administrator account to enable the data collection programs to run correctly. This change is not required for Windows Server 2003.

Procedure

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the Inventory panel, and then select the IBM N series tab.
3. In the Monitoring and Host Configuration capability, select the Data Collection panel.
4. Select the component you want to collect data from.
5. Select or enter the hostname or IP address of the component, and enter a User name and Password with root or administrator rights.
6. Clear the **Save File Locally** check box if you do not want the file copied to your local workstation.
7. Click **Submit**. The diagnostic data is collected in a .tar.gz file.
8. When the Download file dialog box is displayed, select a location to save the file and then click **Save**. This dialog box is displayed only if the **Save File Locally** check box is selected.

Results

The .tar.gz file is stored on the Windows server running VSC for VMware vSphere, and, optionally, copied to your local workstation. On the VSC for VMware vSphere server, the file is saved to the C:\Program Files\IBM\Virtual Storage Console\etc\vsc\web\support.

What to do next

Send the file to Technical Support for analysis.

Changing the service account for data collection on Windows Server 2008

To enable running the data collection programs on Windows Server 2008 or Server 2008 R2, you must change the account used by the Virtual Storage Console for VMware vSphere service to an Administrator account.

About this task

The Virtual Storage Console for VMware vSphere service normally runs under the Local System account. For Windows Server 2008 and Server 2008 R2, the Windows security features prevent the data collection programs from running correctly under the Local System account.

Procedure

1. Log on to the Windows system on which you installed VSC for VMware vSphere.
2. Open the Services application by selecting **Start > Administrative Tools > Services**
3. Right-click Virtual Storage Console for VMware vSphere service and select **Properties**.
4. On the Log On tab, select **This account**.
5. Enter the credentials for an Administrator account on the Windows system and then click **OK**.

What to do next

Switch back to the Local System account if desired after running the data collection programs.

Downloading and running tools

The Monitoring and Host Configuration capability includes tools for detecting and correcting misaligned disk partitions and for setting virtual machine timeouts.

Note: The tools provided by the Monitoring and Host Configuration capability can only be used when the virtual machine (VM) is powered off. The Optimization and Migration capability of VSC allows you to perform online alignments on VMFS-based datastores without having to take your VM down. This capability also lets you review the alignment status of VMs and migrate groups of VMs.

The MBR (master boot record) tools enable you to detect and correct misaligned disk partitions for guest operating systems. You must download these tools before you can use them. There is a set of tools for ESX hosts and one for ESXi hosts. You must download the correct tool set for your hosts.

Enabling the ESXi secure shell

When you are using ESXi, it is a good practice to enable the Secure Shell (SSH) protocol before you download the MBR tools. That way you can use the **scp** command if you need to copy the files. ESXi does not enable this shell by default.

Procedure

1. From an ESXi host, press the key combination ALT F2 to access the Direct Console User Interface (DCUI) screen.
2. Press the F2 function key to get to the Customize System screen.
3. Go to **Troubleshooting Options**.
4. Press **Enter** at the Enable SSH prompt.
5. Press **Enter** at the Modify ESX Shell timeout prompt.
6. Disable the timeout by setting the value to zero (0) and pressing **Enter**.
7. Go to **Restart Management Agents** and press **Enter**.

8. Press F11.

Downloading and installing MBR tools for ESXi hosts

If you have an ESXi host, you must download and install the version of the MBR (master boot record) tools for ESXi. The MBR tools enable you to detect and correct misaligned disk partitions for guest operating systems. These tools must be installed and run directly on the ESXi host. Before you install them, you must extract them from the .tar file into the root directory on the ESXi host.

Before you begin

You must be able to open a console connection to the ESXi host.

Note: The MBR tools can only be used when the virtual machine (VM) is powered off. If you want to perform online alignments on VMFS-based datastores without having to take your VM down, you can use the Optimization and Migration capability. In that case, you do not need to download the MBR tools.

Procedure

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the **Inventory** panel, and then select the IBM N series tab.
3. In the Monitoring and Host Configuration capability, select the Tools panel.
4. Under **MBR Tools**, click the **Download (For ESXi 4.x and ESXi 5.0)** button.
Make sure you download the MBR Tools for ESXi. If you download the wrong MBR tools file, the tools will **not** work.
5. When the File Download dialog is displayed, click **Save**.
6. **(ESXi 4.x)** If you are using ESXi 4.x, manually enable the ESXi shell and SSH so that you can use the **scp** command to copy the files to the correct directories if needed. ESXi 4.x does not enable the ESXi shell and SSH by default. You can enable these options from the physical host or from the vCenter. The following steps enable these options from the vCenter.

Note: vCenter creates a configuration alert for each ESXi host that has the options enabled.

To enable the ESXi shell, perform the following steps:

- a. From vCenter, highlight the appropriate ESXi host.
- b. Go to the **Configuration Tab**.
- c. In the left pane under **Software**, select **Security Profile**.
- d. Select **Properties** from the **Services** pane.
- e. Highlight the **ESXi Shell** service and select **Options**.
- f. Select **Start and Stop with Host**.
- g. Click **Start**.

To enable the ESXi SSH, perform the following steps:

- a. From vCenter, highlight the appropriate ESXi host.
- b. Go to the **Configuration Tab**.
- c. In the left pane under **Software**, select **Security Profile**.
- d. Select **Properties** from the **Services** pane.
- e. Highlight the **SSH** service and select **Options**.
- f. Select **Start and Stop with Host**.

g. Click **Start**.

7. Copy the MBR tools for ESXi file to the root (/) directory of the ESXi host. If you are using ESXi 4.x, use the Troubleshooting Console. If you are using ESXi 5.x, use the Technical Service Console. You might need to open ESXi firewall ports to enable copying the tools to the host.

Note: The MBR tools libraries must be located in specific directories on the host. Be sure to download the file to the root directory of the ESXi host.

8. Extract the files by entering the following command: `tar -zxvf mbrtools_esxi.tgz`. If you did not download the file to the root directory, you must manually move the files to that directory.

Note: ESXi does not support `-P` with the `tar` command.

What to do next

Run the `mbralign` tool to check and fix the partition alignment.

Downloading and installing MBR tools for ESX hosts

If you have an ESX host, you must download and install the version of the MBR tools for ESX. These tools must be installed and run directly on the ESX host. You cannot run these tools from the vSphere Client, vCenter Server, or VSC for VMware vSphere server.

Before you begin

You must be able to open a console connection to the ESX host.

Procedure

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the **Inventory** panel, and then select the IBM N series tab.
3. In the Monitoring and Host Configuration capability, select the Tools panel.
4. Under **MBR Tools**, click the appropriate **Download** button that states for **(For ESX 4.x)**.

If you download the wrong MBR tools file for your host, the tools will **not** work.

5. When the File Download dialog is displayed, click **Save**.
6. Copy the MBR tools file to the root (/) directory of the host computer. You might need to open ESX firewall ports to enable copying the tools to the host.

Note: The MBR tools libraries must be located in specific directories on the host. Be sure to download the file to the root directory.

7. Extract the files on the ESX host using the by entering the following command:
`tar -Pzxf mbrtools.tar.gz`

The `-P` option places the files in the required directories. You can move the binary files to any location, but the library files must be located in the specific directories to which they are originally extracted.

What to do next

Run the `mbralign` tool to check and fix the partition alignment.

Related concepts:

“How to identify and fix VMDK partition alignment issues” on page 178

“Optimizing and migrating datastores and virtual machines” on page 87

Installing GOS scripts

The ISO images of the guest operating system (GOS) scripts are loaded on the VSC for VMware vSphere server. Mount and run them from the vSphere Client to set the storage timeouts for virtual machines.

Before you begin

The virtual machine must be running.

The CD-ROM must already exist in the virtual machine or it must be added.

The script must be installed from the copy of the VSC for VMware vSphere registered to the vCenter Server that manages the VM.

Procedure

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the **Inventory** panel, and then select the IBM N series tab.
3. In the Monitoring and Host Configuration capability, select the Tools panel.
4. Under **Guest OS Tools**, right-click the link to the ISO image for your guest operating system version and select **Copy to clipboard**.
5. In the vSphere Client, select the desired VM and click the **CD/DVD Connections** icon.
6. Select **CD/DVD Drive 1 > Connect to ISO image on local disk**.
7. Paste the link you copied into the **File Name** field and then click **Open**. If you receive an authorization error, be sure you select the IBM N series tab and click **Yes** to proceed if a security certificate warning is displayed.

Also, be sure that the link you are using is from the copy of the VSC for VMware vSphere running on the vCenter Server that manages the VM.

What to do next

Log on to the VM and run the script to set the storage timeout values.

Related tasks:

“Running the GOS timeout scripts for Linux” on page 176

“Running the GOS timeout scripts for Solaris” on page 177

“Running the GOS timeout script for Windows” on page 177

Adding the CD-ROM to a VM:

Add the CD-ROM to a virtual machine if it does not exist to enable installing the guest operating system scripts.

Procedure

1. In the vSphere Client, select the desired VM and power it off.
2. Right-click the virtual machine and select **Edit Settings**.
3. On the **Hardware** tab, click **Add**.
4. Select **CD/DVD Drive** and then click **Next**.

5. Click **Use physical drive**.
6. Click **Next** several times to accept the default values.
7. Click **OK** to finish adding the CD-ROM.
8. Power on the VM.

Fields and commands described by the online help

In addition to the tasks that this guide has already described, the online help contains explanations for other fields and commands that are available on the Monitoring and Host Configuration capability panels. To provide you with an overview of those additional fields and commands, they are included in the sections that follow.

If you are already familiar with the Monitoring and Host Configuration capability panels, you can skip these sections.

Overview panel fields and commands

The Overview panel displays general information about storage controllers and ESX and ESXi hosts that are discovered by VSC for VMware vSphere.

Navigation tips

This panel contains a section for storage controllers and another section for hosts. In the default view, it shows several columns for each section. You can adjust the display to show information that is most relevant to you as well as add or remove columns. There are several ways to change the display:

- Click an arrow next to a column heading to display a drop-down menu that lets you specify the sort method, the displayed columns, and other options.
You can modify the current column display by selecting the **Columns** option. In the pop-up box that appears, select the column names.
- Right-click on a row to display a pop-up menu. In the Storage controller section, this menu contains several options, including opening Element Manager or FilerView, viewing cluster LIF details, viewing connected hosts, modifying controller credentials, and adding or removing a controller. In the ESX Hosts section, this menu provides options for setting recommended values, showing details about the host, and skipping the host.
- Use the scroll bar to move through the controller or host section.

Adding, removing, modifying storage systems

This panel provides two ways to add or remove a storage system (also called "controller") or modify the credentials of a storage system.

- Select the **Add**, **Delete**, or **Edit** buttons in the upper right of the panel.
- Right-click on a storage system or host to display a pop-up menu with these options.

Both methods display the same pop-up dialog boxes. These dialog boxes are explained below in the sections **Add Storage System**, **Remove Controller**, and **Modify Credentials**.

When you access VSC for VMware vSphere from the Inventory section of the vSphere Client, the Monitoring and Host Configuration panels display only the resources associated with the selected object.

Update button

The **Update** button starts the discovery process and updates VSC for VMware vSphere with any added storage controllers or ESX or ESXi hosts. After an update, all clients connected to the same plug-in are automatically refreshed with the new configuration data.

Note: Only those resources associated with the object selected in the Inventory panel are updated. To update all resources, select the Datacenter object in the Inventory panel before clicking **Update**. If the Inventory panel is not displayed because you accessed VSC for VMware vSphere icon in the **Solutions and Applications** section of the vSphere Client **Home** page, then all resources are updated.

Storage Controllers

The **Storage Controllers** section of this panel displays several columns of information about the storage controller.

Controller

This column displays the hostname reported by the controller. It is divided into sections based on the controller category. Each category contains listings for all the controllers in that category. For example, suppose the category is "HA Pair:vmshe-06/vmshe-07". There would be one row for Controller:vmshe-06 and another row for Controller:vmshe-07. VSC for VMware vSphere uses the following identifiers for controllers:

- **Unknown**

A controller that has not been discovered.

- **Controller**

A stand-alone controller. The hostname has the prefix "controller".

- **HA Pair**

One or more controllers. In this case the controller is running a version of Data ONTAP that is **not** operating in Cluster-Mode. The hostname has the prefix "controller".

- **MultiStore**

A controller for a MultiStore vFiler unit. The hostname has the prefix "MultiStore".

Note: By default, the Monitoring and Host Configuration capability does not show unassociated vFiler units. If you manually add a vFiler unit, it is displayed only if you are at the Datacenter level or higher.

- **Cluster**

A controller pair running Data ONTAP operating in Cluster-Mode. Based on the controller, the hostname can have one of the following prefixes:

- **Cluster**
- **Node**
- **Vserver**

If you have the correct credentials, VSC for VMware vSphere discovers the entire cluster. If you only have credentials for the Vserver management LIF and do not have the correct credentials to connect to the cluster-management LIF, VSC for VMware vSphere

displays only the SAN and NAS datastores associated with the Vserver. No cluster names, nodes or other Vservers on the cluster are displayed.

If you connect the cluster directly to the Vserver, not all the VSC for VMware vSphere features are supported. For example, VSC for VMware vSphere cannot perform an upfront validation of the RBAC privileges even though RBAC is fully supported. The upfront validation occurs only when the controller attaches to a cluster-management LIF.

Note: By default, the Monitoring and Host Configuration capability does not show Vservers that are not attached to an ESX/ESXi host that is being managed by a vCenter Server. If you manually add a Vserver, it is displayed only if you are at the Datacenter level or higher to see it. If the Vserver does not have connected datastores, the Monitoring and Host Configuration capability does not display it even if you manually add it.

IP Address

The IP address that maps to the hostname.

If you have configured Vserver management LIFs, VSC for VMware vSphere returns IP addresses for those LIFs as well as for the IP addresses associated with the cluster-management LIF. If the correct credentials are not associated with the IP address, VSC for VMware vSphere is not able to display this information. In that case you must select the IP address and provide credentials.

Version

The version of Data ONTAP software that the controller is running. If the controller is running Data ONTAP operating in Cluster-Mode, the version has "Cluster-Mode" appended to the end.

Status

The overall status reported by the controller.

Note: You can display the **Status Reason** column to get more information about why the controller has the status shown. To display this column, click the arrow in the column heading to display the pop-up box and select **Columns**. In the pop-up box that appears, select the check box next to Status Reason.

Status values are:

- **Normal** - No problems reported.
- **Alert** - Need to investigate; more information in the **Status Reason** column.

For vFiler units, the **Alert** icon means that the vFiler unit does not respond to VSC for VMware vSphere. No status reason is reported. Connect to the physical controller that owns the vFiler unit to get more information.

- **Insufficient Privileges** - One or more required privileges are not specified for the custom user name provided for this controller. Specify a different user name or update RBAC roles. For information about configuring RBAC,

Note: If the storage is directly connected to a Vserver, upfront validation of the RBAC privileges is not done. This check is only performed when the storage controller connects to a cluster-management LIF.

- **Authentication Failure** - At initial discovery, unable to log into the controller using the default credentials.

At any other time, unable to log into the controller using the specified credentials. Right-click the controller, select **Modify Credentials**, and then enter a valid user name and password.

Instead of entering credentials for each controller, you can specify new default credentials. To do this, click **Set Default Controller Credentials** on the Discovery panel.

Note: The Monitoring and Host Configuration manages the controller credentials for all the capabilities except Backup and Recovery.

- **Skipped** - This controller has been designated as skipped so discovery is not attempted.

To use this controller, right-click the controller name and select **Modify Credentials**. Enter new credentials, if needed, and clear the **Skipped** check box.

Note: If a controller is skipped, the Monitoring and Host Configuration capability does not export it to the Optimization and Migration and Provisioning and Cloning capabilities. If it exists for these capabilities, it will be deleted from their controller lists and unavailable for their work flows. However, if you clear the **Skipped** check box, the Monitoring and Host Configuration capability adds the controller back to the export list and the Optimization and Migration and Provisioning and Cloning capabilities can add it to the work flows again.

- **SSL is not configured** - The controller is not configured to use a secure connection. Enable SSL on the controller, if possible. Otherwise right-click the controller, select **Modify Credentials**, and clear the **Use SSL** check box.
- **Unmanaged** - The version of Data ONTAP software on the controller does not support the SAN functionality of VSC for VMware vSphere. The **Storage Details - SAN** panel in VSC for VMware vSphere does not display any details for unmanaged controllers. The NAS functionality of VSC for VMware vSphere works with some earlier versions of Data ONTAP that lack the interfaces needed for SAN support. See the Interoperability Matrix for the required version of Data ONTAP software.
- **Unknown** - This controller has not been discovered.
- **Updating** - The update task is in progress.

Free Capacity

This value is specific to each controller. It is the controller's total disk capacity minus the total used capacity on that controller. VSC for VMware vSphere reports space usage and capacity information from varying perspectives. In this situation, the total capacity for a controller is the sum of the raw sizes of all the disks that are not spares, partner disks, or parity disks. The used space is the total space that is used in all the aggregates on the controller.

For Vservers and vFilers units, the Monitoring and Host Configuration capability calculates the values for free space, used space, and total space based on the free space, used space, and total space of all the volumes owned by that vServer or vFiler unit.

VAAI Capable

Displays whether the storage controller is running a version of Data ONTAP software that supports the VMware vStorage APIs for Array Integration (VAAI) features. This feature uses storage system intelligent services to create copies of VMDKs without passing data through the ESX or ESXi host. Values are:

- **Supported** - The controller supports VAAI and the VAAI support is enabled for use. Controllers running Data ONTAP operating in Cluster-Mode use this designation.
- **Enabled** - The controller supports VAAI and the VAAI support is enabled for use.
Controllers running Data ONTAP operating in 7-Mode use this designation.
- **Disabled** - The controller is capable of supporting VAAI, but the VAAI support is disabled on the controller.
This designation applies to all controllers.
- **No** - The controller is not running a version of Data ONTAP that is capable of supporting VAAI.

Supported Protocols

The storage protocols licensed on the controller. Options are FC, NFS, iSCSI, or None. Does not imply the protocol is in use by an ESX or ESXi host. If there is a problem, this value is displayed as Unknown.

For MultiStore vFiler units, this field reports the protocols actually in use by ESX or ESXi hosts. Options are NFS and iSCSI.

For Nodes, this field is reported as N/A.

You can display the following fields by clicking the arrow in the column header and selecting **Columns**.

Partner

The hostname of the partner controller in an active/active storage configuration.

Status Reason

The detailed storage controller status returned by the controller. Note that vFiler units do not return a detailed status reason. Instead, This controller is a MultiStore vFiler unit is displayed for vFiler units. The displayed text is truncated if the status reason is too long, and ellipses (...) are appended to the displayed text.

Right-click on the row to display a pop-up box that lets you use either Element Manager or FilerView to connect to the storage controller and view more status information. The option that appears depends on which mode of Data ONTAP the storage system is running.

Total Used

The sum of all the space reported as used by all the aggregates on the controller. If the controller is either a Multistore vFiler Unit or a Vserver

and is running Data ONTAP operating in Cluster-Mode, this is the sum of all the used space in all the volumes owned by either the vFiler or the Vserver.

Note: For Vservers and vFiler units, the Monitoring and Host Configuration capability calculates the used space based on the used space of all the volumes owned by that vServer or vFiler unit.

Total Allocated

The total space currently allocated in all the aggregates on the controller. If the controller is either a Multistore vFiler Unit or a Vserver and is running Data ONTAP operating in Cluster-Mode, this is the sum of all the space currently allocated in all the volumes owned by either the vFiler or the Vserver.

Total Capacity

The total physical disk space on the controller. If the controller is either a Multistore vFiler Unit or a Vserver and is running Data ONTAP operating in Cluster-Mode, this is the sum of all the space currently allocated in all the volumes owned by either the vFiler or the Vserver.

Note: For Vservers and vFiler units, the Monitoring and Host Configuration capability calculates the total space based on the total space of all the volumes owned by that vServer or vFiler unit.

CF Mode

The cfmode setting of the controller, if it is available.

CF Status

The failover status of the controller in an active/active configuration. The normal status is reported as "CONNECTED". During a takeover, the controller that takes over for its partner reports "TAKEOVER", and the partner controller that has been taken over reports "TAKEN OVER". If failover is not enabled, the field is blank.

Different commands are available depending on the status of the selected controller. The right-click menu for each entry in this section has the following commands:

Add Storage System

Manually add a new storage system to be managed by VSC for VMware vSphere.

The default management port number is 443 if the SSL check box is selected and 80 if it is not selected. These are the Data ONTAP defaults. If you toggle the SSL check box, the port number switches between 443 and 80. You can specify a different port number. If you do that, then toggling the SSL check box only changes the SSL state in the dialog box.

Note: If the IP address you supply when you add a storage system connects to Multistore vFiler unit, you should change it to an IP address that connects to a physical storage controller unit in order to take advantage of all the functions available.

Remove Controller

Removes a storage system from VSC for VMware vSphere. If the storage system has LUNs or NFS exports mapped to an ESX or ESXi host managed by VSC for VMware vSphere, the storage system cannot be removed, and an error message is displayed.

Modify Credentials

Enables you to set certain information associated with the login credentials for the selected storage system.

Note: The Monitoring and Host Configuration capability manages credentials for the Optimization and Migration capability and Provisioning and Cloning capability. You cannot enter separate credentials for those two capabilities.

Credentials are assigned based on the user name/password pair. You cannot change the roles associated with that user name/password pair at the Modify Credentials dialog box. You can specify the following:

- Management IP address
- Management port number
- Whether SSL is enabled
- Whether the host is skipped

VSC for VMware vSphere uses the management IP address to communicate with the storage system. VSC for VMware vSphere lists the available addresses.

The default management port number is 443 if the SSL check box is checked and 80 if it is not checked. These are the Data ONTAP defaults. If you toggle the SSL check box, the port number switches between 443 and 80. You can specify a different port number. If you do that, then toggling the SSL check box only changes the SSL state in the dialog box.

If you chose to **not** provide credentials for this storage system, the **Skipped** check box is selected. Clear the check box and enter the credentials for the storage system.

After you have supplied the correct information for this dialog box, you can click **OK** to see the list of allowed and disallowed roles for this user name/password pair. Clicking **OK** again assigns those values to the controller. Clicking **Cancel** returns you to the dialog box where you can enter a different user name/password that is associated with different roles.

Note: Instead of entering credentials for each storage system, you can specify new default credentials. Go to the Discovery panel and click **Set Default Controller Credentials**.

Open Element Manager

Opens a browser window and connects to Element Manager for the selected storage controller. Prompts for login credentials if needed. This feature works with controllers running Data ONTAP operating in Cluster-Mode.

Note: Element Manager is not available for storage systems running Data ONTAP operating in 7-Mode.

Open FilerView

Opens a browser window and connects to the FilerView GUI for the selected storage controller. Prompts for login credentials if needed. This feature works with controllers running Data ONTAP operating in 7-Mode.

Note: FilerView is not available for vFiler units or storage systems running Data ONTAP operating in Cluster-Mode. If you are using Data ONTAP

operating in Cluster-Mode, you can download a copy of OnCommand System Manager and use that tool for managing LUNs.

View Connected Hosts

Displays a list of ESX or ESXi hosts using storage on this controller.

ESX Hosts

The **ESX Hosts** section of this panel contains the following fields that apply to both ESX and ESXi hosts. Click the column heading to get sort options and to modify the columns displayed.

Hostname

The hostname of the ESX or ESXi host.

IP Address

The IP address of the ESX or ESXi host that maps to the hostname.

Version

The version of ESX or ESXi server running on the host. This numerical version does not include updates. For example, ESX 4.0U1 is reported as 4.0.0. Displays **Unknown** if VSC for VMware vSphere is unable to communicate with the host.

Status

The overall status of the host.

Note: You can display the **Status Reason** column to get more information about why the host has the status shown. To display this column, click the arrow in the column heading to display the pop-up box and select **Columns**. In the pop-up box that appears, check the box next to Status Reason.

Status values are:

- **Normal** - No problems reported.
- **Alert** - vCenter reported a problem with the overall health of the host. The Alert status is based on all alarms associated with the host. See the vCenter **Alarms** tab for the host to see the alarm or alarms that caused this status.
- **Pending Reboot** - The recommended storage adapter or NFS values have been set by VSC for VMware vSphere, but the host needs to be rebooted to make the settings take effect.

Note: Installing the NFS Plug-in for VMware VAAI also results in a status value of **Pending Reboot**.

- **Skipped** - A user chose to skip this host by right-clicking on the hostname selecting **Skipped**. In this case, discovery is not attempted and best practice settings cannot be applied to this host. The columns for Adapter, MPIO, and NFS settings display as **Unknown**. See the **Status Reason** column for an explanation.

To use this host, right-click the hostname and clear the Skip host check box.

- **Unknown** - VSC for VMware vSphere is unable to communicate with the host.
- **Updating** - The update task is in progress.

Adapter Settings

Displays whether the recommended storage adapter settings are applied to this host. You can right-click the host and select **Set Recommended Values** to update the adapters. Values are:

- **Normal** - Recommended adapter settings applied.
- **Alert** - Need to set recommended settings. Right-click the host and select **Show Details** to view current settings.
- **Unknown** - VSC for VMware vSphere is unable to communicate with the host.

MPIO Settings

Displays whether the recommended multipath I/O settings are applied to this host. You can right-click the host and select **Set Recommended Values** to update the host's MPIO settings. Values are:

- **Normal** - Recommended MPIO settings applied.
- **Alert** - Need to set recommended settings. Right-click the host and select **Show Details** to view current settings.

If a storage controller is skipped during discovery, VSC for VMware vSphere is unable to set the MPIO settings for the paths to that controller. An error message is displayed.

- **Unknown** - VSC for VMware vSphere is unable to communicate with the host.

NFS Settings

Displays whether the recommended NFS settings are applied to this host. You can right-click the host and select **Set Recommended Values** to update the host's NFS settings. Values are:

- **Normal** - Recommended NFS settings applied.
- **Alert** - Need to set recommended settings. Right-click the host and select **Show Details** to view current settings.
- **Unknown** - VSC for VMware vSphere is unable to communicate with the host.

You can display the following field by clicking the arrow in the column header and selecting **Columns**.

Status Reason

The reason for the host's current status, such as "Pending Reboot" or "Skipped". If the host is throwing an alert, this field directs the user to the **Alert** tab for this host. The displayed text is truncated if the status reason is too long, and ellipses (...) are appended to the displayed text.

The right-click menu for each entry in this section has the following commands.

Set Recommended Values

Sets the recommended values for storage adapters, multipath I/O, and NFS. A host reboot is required after setting adapter and NFS settings.

Show Details

Displays detailed information about the host, including storage adapter driver versions and settings, NFS mounts, and mapped LUNs.

Skip Host

Lets you tell VSC for VMware vSphere to skip this host and not attempt to discover it.

Storage Details - SAN panel fields and commands

The panel for Storage Details - SAN displays detailed information about storage controllers and LUNs used by ESX or ESXi hosts managed by VSC for VMware vSphere.

SAN Storage Controllers

The **Storage Controllers** section of this panel contains numerous columns. You can modify the column display by clicking the arrow next to a column heading and selecting the **Columns ...** option from the drop-down menu that appears. Check or uncheck the column names in the pop-up box that appears to specify which columns VSC for VMware vSphere displays.

Note: When you access VSC for VMware vSphere from the Inventory section of the vSphere Client, the Monitoring and Host Configuration panels display only the resources associated with the object selected in the Inventory panel. This enables you to focus on only the resources you are currently interested in.

Datastore

The name of the VMFS datastore on the LUN reported by the ESX or ESXi host. For a raw device mapping (RDM) LUN, this field displays -- RDM LUN --. If the datastore is comprised of extents, then this field displays the name as *datastore_name-extent*.

Datastore Capacity

The total capacity of the VMFS datastore reported by the ESX or ESXi host. Displays "N/A" for RDM LUNs.

LUN Pathname

The path to the LUN on the storage controller.

Storage Capacity

The total LUN size reported by the storage controller.

Storage Status

The LUN status reported by the storage controller. Values are:

- **Online** - The LUN is available.
- **Offline** - The LUN is not available.

Thin Prov

The thin provisioning status of the LUN reported by the storage controller. Values are:

- **Disabled** - Physical storage for the entire LUN is reserved on the storage controller.
- **Enabled** - Physical storage for the entire LUN size reported to the ESX or ESXi host is not reserved on the storage controller. Physical storage is allocated to the LUN as needed. The volume containing the LUN must be managed to prevent running out of space.

You can display the following fields by clicking the arrow in the column header and selecting **Columns**.

Controller

The hostname of the controller that owns the LUN.

Partner

The hostname of the partner controller in an active/active storage configuration.

LUN ID

The identifier (number) to which the LUN is mapped.

Serial Number

The serial number of the LUN on the controller.

Volume

The volume on the storage controller that contains the LUN.

Aggregate

The aggregate on the storage controller that contains the LUN.

Displays "N/A" for vFiler units. Direct vFiler units do not have aggregates to display.

Details

The **Details** section of this panel contains the following fields.

LUN Displays detailed information about the selected LUN including whether space reservations are enabled, LUN type, protocol, igroup name, and whether ALUA is enabled on the igroup.

Deduplication (Advanced Single Instance Storage)

Displays whether deduplication is enabled for the LUN, and if so, how much storage is recovered, the date and time of the last deduplication, and the deduplication schedule.

Note: If VSC for VMware vSphere is communicating directly with a Vserver, the deduplication fields are left blank. This is because the default vsadmin role does not have sufficient privileges to access this information.

Capacity

Displays the percentage of space used in the datastore, LUN, volume, and aggregate.

Place your cursor (hover) over these fields to display the details. For example, place your cursor over the **Volume Usage** field to display a breakdown of total volume size into data and Snapshot reserve used and free.

Volume

Displays the name of the volume that contains the LUN and detailed volume settings.

Actions

The following action is available on the Storage Details - SAN panel:

View Mapped Hosts

Select a LUN under **SAN Storage Controllers**, and then click this link to display which ESX or ESXi hosts access the LUN. This link is located under the Capacity detail box.

Storage Details - NAS panel fields and commands

The Storage Details - NAS panel displays detailed information about storage controllers and NFS exports (datastores) used by ESX or ESXi hosts managed by VSC for VMware vSphere.

NAS Storage Controllers

The **NAS Storage Controllers** section of this panel contains numerous columns. You can modify the column display by clicking the arrow next to a column heading and selecting the **Columns ...** option from the drop-down menu that appears. Check or uncheck the column names in the pop-up box that appears to specify which columns VSC for VMware vSphere displays.

Note: When you access VSC for VMware vSphere from the Inventory section of the vSphere Client, the Monitoring and Host Configuration panels display only those resources associated with the object selected in the Inventory panel. This enables you to focus on only those resources currently of interest.

Datastore

The name of the NFS datastore reported by the ESX or ESXi host. For aliased NFS mounts, this is the alias name.

Datastore Capacity

The total capacity of the NFS datastore reported by the ESX or ESXi host.

NFS Pathname

For storage controllers running Data ONTAP operating in Cluster-Mode, this is the junction path to the mounted NFS export.

For storage controllers running Data ONTAP operating in 7-Mode, this is the path to the qtree or volume on the controller for the NFS export.

Data Path Access

Whether a path setting is Direct (green check), Indirect (red exclamation point (!)), N/A, or (unknown).

With Data ONTAP operating in Cluster-Mode, it is possible for a client to access a data LIF using an indirect data path to a FlexVol. This happens when a remote data LIF is bound to a different physical node than the one that owns the exported FlexVol. To have a direct data path, the client must access a data LIF that is local to the node that owns the exported FlexVol.

If you have an indirect data path setting, you can right-click that row and select the **View Direct Data Path Choices** option. This option displays the **Direct Data Path Choices** pop-up box, which contains a list of ports to direct data paths. You cannot move the LIF to the direct data path from this pop-up box. That needs to be done by a storage system administrator who has the correct credentials and is using either the storage system console or a tool such as System Manager.

You can, however, decide which port you want to use. Then you must make sure that port is connected to the network. VSC for VMware vSphere does not check the network connectivity between the port and the ESX/ESXi host. While most of the ports are connected to the network, if you use one that is not, your datastore will go offline when the path is switched.

After you know the port is connected to the network, collect the information that is displayed in the **Direct Data Path Choices** pop-up and give it to the storage system administrator. While you must manually verify that the port is connected to the network, the **Direct Data Path Choices** pop-up contains all the other information a storage administrator needs to move the LIF

If multiple datastores are using the LIF, the data path access for those other datastores will change from direct to indirect after the LIF is moved.

If a direct Vserver connection is made, VSC for VMware vSphere cannot query the storage controller to determine the path.

An N/A path indicates a path to a storage controller running Data ONTAP operating in 7-Mode.

An unknown path occurs if the discovery data is incomplete.

Storage Status

The status of the qtree or volume for the NFS export reported by the controller. Values are:

- **Normal** - The qtree or volume is available.
- **Alert** - The qtree or volume is not available.

You can display the following fields by clicking the arrow in the column header and selecting **Columns**.

Controller

The hostname of the controller that owns the datastore (NFS export).

Partner

The hostname of the partner controller in an active/active storage configuration.

Volume

The volume on the storage controller that contains the NFS export or VSC for VMware vSphere is communicating directly with a Vserver.

Aggregate

The aggregate on the storage controller that contains the NFS export.

Displays "N/A" for vFiler units, which do not have their own aggregates.

Details

The **Details** section of this panel contains the following fields.

NFS Details of the NFS export, including its pathname on both the storage controller and ESX or ESXi host, the file system security style, and the user name granting root access to the directory.

For aliased NFS mounts, the actual NFS path name is shown in the **NFS pathname** field, with "(actual)" added after the name.

Host Privileges

Click these links to list the hosts with read, read/write, and root access.

Deduplication (Advanced Single Instance Storage)

Displays whether deduplication is enabled for the LUN, and if so, how much storage is recovered, the date and time of the last deduplication, and the deduplication schedule.

Note: If VSC for VMware vSphere is communicating directly with a Vserver, the deduplication fields are left blank. This is because the default vsadmin role does not have sufficient privileges to access this information.

Capacity

Displays the percentage of space used in the datastore, volume, and aggregate.

Place your cursor (hover) over these fields to display the details. For example, place your cursor over the **Volume Usage** field to display a breakdown of total volume size into data and Snapshot reserve used and free.

Volume

Displays the name of the volume that contains the datastore and detailed volume settings.

Actions

The following actions are available on the Storage Details - NAS panel:

View Read-Only Hosts (Host Privileges)

Select a datastore under **NAS Storage Controllers** and then click this link to display which ESX or ESXi hosts have read-only permission for the datastore.

View Read-Write Hosts (Host Privileges)

Select a datastore under **NAS Storage Controllers** and then click this link to display which ESX or ESXi hosts have read-write permission for the datastore.

View Root Access Hosts (Host Privileges)

Select a datastore under **NAS Storage Controllers** and then click this link to display which ESX or ESXi hosts have root access permission for the datastore.

View Mounted Hosts

Select a datastore under **NAS Storage Controllers** and then click this link to display which ESX or ESXi hosts mount the selected NFS export.

Data Collection panel fields and commands

The Data Collection panel enables you to collect troubleshooting information from storage controllers, FC switches, and ESX or ESXi hosts.

The data you collect should be sent to Technical Support personnel for analysis.

The fields on this panel are as follows:

Target Hostname

The hostname or IP address of the device you are collecting data from. For storage controllers and ESX or ESXi hosts, you can select the name from the pull-down menu.

Note: IPv6 addresses are not currently supported.

User Name

The user name you provide must have Administrator or root privileges on the device. For storage controllers, VSC for VMware vSphere fills in the stored credentials.

Password

The password for the user name. For storage controllers, VSC for VMware vSphere fills in the stored credentials.

Collect information from

The type of device. Note that the switch options are for supported FC switches.

Export VSC Logs

This option collects the log files from all installed VSC for VMware vSphere components and capabilities.

Save File Locally

Specifies that a copy of the data collected is saved to your local workstation. If you select this option, VSC prompts you for a path where it will save the file after the data is collected. The file is always saved on the server running VSC for VMware vSphere.

The data is saved in a .tar.gz file that is stored on the Windows server running VSC for VMware vSphere and, optionally, copied to your local workstation. On the server, the file is saved to the \webapps\public\support folder in the VSC for VMware vSphere installation directory. The default path is C:\Program Files\IBM\Virtual Storage Console\webapps\public\support\.

Note: The data collection can take a few minutes. While the data is being collected, VSC for VMware vSphere might not respond to user input.

Tools panel fields and commands

The Tools panel provides software tools you can run on individual ESX or ESXi hosts.

MBR Tools

The master boot record (MBR) tools enable you to detect and correct misaligned disk partitions for guest operating systems. To use these tools, you must power off the virtual machine (VM).

Note: To check the alignment status of VMs, perform online alignments of SAN-based datastores, and migrate groups of VMs, use the Optimization and Migration capability.

The tools must be installed and run directly on the ESX or ESXi host. They cannot run from the vSphere Client, vCenter Server, or VSC for VMware vSphere server.

There are two **Download** buttons: one for ESX hosts and one for ESXi hosts. You must click the correct **Download** button for your host. The **Download** button opens the File Download dialog and enables you to save a copy of the tools software package.

The ESX MBR tools support a copy offload feature for NFS datastores. The ESXi MBR tools **do not** support this feature.

Note: The MBR tool libraries must be located in specific directories on the ESX or ESXi host. If you are using ESXi, be sure to download the tools library file to the root directory of the host.

Guest OS Tools

The guest OS timeout scripts set the storage timeout values for supported Linux, Solaris, and Windows guest operating systems. The timeout values ensure correct failover behavior.

You can mount and run the scripts from the vSphere client to set the storage timeouts for virtual machines. You can mount the URL for the appropriate

operating system as a virtual CD-ROM in the virtual machine using the vSphere client. You then run the script from the virtual machine's console.

You can right-click the tool's URL to copy it to the clipboard.

Note: The guest OS scripts must be installed from the copy of VSC for VMware vSphere registered to the vCenter Server that manages the VM.

NFS Plug-in for VMware VAAI

The NFS Plug-in for VMware VAAI is a software library that integrates with the VMware Virtual Disk Libraries. When you install these libraries on ESXi 5.0 hosts, VMware can execute various primitives on files stored on the storage systems.

Note: The NFS Plug-in for VMware VAAI is supported with vSphere 5.0 ESXi hosts and certain versions of Data ONTAP: Data ONTAP 8.1 or later when operating in Cluster-Mode and Data ONTAP 8.1.1 or later when operating in 7-Mode. The plug-in is not supported with earlier versions of ESXi or Data ONTAP.

This plug-in is not shipped with VSC for VMware vSphere.

After you download the NFS Plug-in for VMware VAAI and place it in the correct directory, click the Monitoring and Host Configuration capability **Install on Host** button. Selecting this button starts an automatic installation process. VSC for VMware vSphere displays a pop-up dialog box that allows you to specify the ESXi hosts on which you want to install the plug-in.

Note: The Monitoring and Host Configuration capability checks the versions of ESXi on the host and Data ONTAP on the storage systems attached to the host. If the versions are not correct, the dialog box grays out the name of that host. You can only select hosts that have the correct version of ESXi and vSphere and are connected to storage systems running the correct version of Data ONTAP.

VSC for VMware vSphere displays the progress of the installation. After the plug-in is installed, you must restart the host to start the NFS Plug-in for VMware VAAI.

Check the *Release Notes* for the most current information about the NFS Plug-in for VMware VAAI.

Discovery Status panel fields

The Discovery Status panel displays summary information about VSC for VMware vSphere and the number of resources it has discovered. You can also set default storage controller credentials for discovery from the panel.

Selected Object Totals Fields

The fields in the first column on this panel display the totals for resources associated with the object selected in the Inventory panel of the vSphere client. If the Inventory panel is not displayed, this column includes the totals for all resources discovered by VSC for VMware vSphere.

Total number of Hosts

The total number of ESX or ESXi hosts managed by VSC for VMware vSphere.

Non-compliant Hosts

The number of ESX or ESXi hosts that do not have the correct adapter, MPIO, or NFS settings. You can use VSC for VMware vSphere to update the settings from the Overview panel.

Total number of Controllers

The total number of storage controllers running Data ONTAP operating in 7-Mode or Data ONTAP 7G software that have been discovered by VSC for VMware vSphere. An active/active storage configuration is counted as two controllers, assuming both controllers provide storage to ESX or ESXi hosts. The partner of a controller in an active/active configuration is not reported or managed if it does not provide storage to an ESX or ESXi host. This total includes Multistore vFiler units.

MultiStore vFiler units

The number of discovered storage controllers that are MultiStore vFiler units.

Total number of Clusters

The total number of storage controller clusters discovered by VSC for VMware vSphere.

Number of Vservers

The number of Vservers in a Data ONTAP cluster.

Number of LUNs

The total number of LUNs used by ESX or ESXi hosts for which the storage controller is known and the controller credentials are supplied.

Number of NFS Mounts

The total number of NFS mounts used by ESX or ESXi hosts for which the storage controller is known and the controller credentials are supplied.

Note: For configurations using vSphere 5.0 or later, the maximum number of mounted NFS volumes has increased to 256 from 64. As a result, VSC for VMware vSphere uses 256 as the default value for `NFS.MaxVolumes`, 32 as the default value for `Net.TcpIpHeapSize`, and 128 as the default value for `Net.TcpIpHeapMax` for these configurations.

vCenter Totals Fields

The fields in the second column on this panel display the totals for all resources discovered by VSC for VMware vSphere.

Total number of Hosts

The total number of ESX or ESXi hosts managed by VSC for VMware vSphere.

Non-compliant Hosts

The number of ESX or ESXi hosts that do not have the correct adapter, MPIO, or NFS settings. You can use VSC for VMware vSphere to update the settings from the Overview panel.

Total number of Controllers

The total number of storage controllers discovered by VSC for VMware vSphere. An active/active storage configuration is counted as two controllers, assuming both controllers provide storage to ESX or ESXi hosts. The partner of a controller in an active/active configuration is not reported or managed if it does not provide storage to an ESX or ESXi host. This total includes Multistore vFiler units.

MultiStore vFiler units

The number of discovered storage controllers that are MultiStore vFiler units.

Total number of Clusters

The total number of storage controller clusters discovered by VSC for VMware vSphere.

Number of Vservers

The number of Vservers.

Number of Nodes

The number of storage controller clusters that are nodes.

Number of LUNs

The total number of LUNs used by ESX or ESXi hosts for which the storage controller is known and the controller credentials are supplied.

Number of NFS Mounts

The total number of NFS mounts used by ESX or ESXi hosts for which the storage controller is known and the controller credentials are supplied.

Note: For configurations using vSphere 5.0, the maximum number of mounted NFS volumes has increased to 256 from 64. As a result, VSC for VMware vSphere uses 256 as the default value for `NFS.MaxVolumes` and 128 as the default value for `Net.TcpIpHeapMax` for these configurations.

Number of Failed Controllers

The total number of storage controllers that cannot be fully managed by VSC for VMware vSphere. This total includes the following failure types:

Authentication Failure

The number of storage controllers for which credentials are not valid. For newly discovered controllers, VSC for VMware vSphere was unable to connect using the default credentials. You can either update the credentials for an individual controller on the Overview panel, or you can set new default credentials by clicking **Set Default Controller Credentials** at the bottom of this panel.

Insufficient Privileges

The number of storage controllers for which the supplied credentials are valid, but lack one or more required privileges associated with the RBAC role. You can either specify a different user name or update the RBAC privileges for the specified custom user name. For information about configuring RBAC, check the N series support website (accessed and navigated as described in Websites).

Note: If the storage is directly connected to a Vserver, upfront validation of the RBAC privileges is not done. This check is only performed when the storage controller connects to a cluster-management LIF.

SSL not configured

The number of storage controllers for which SSL is not configured. Using SSL to communicate with storage controllers is strongly recommended.

Unknown Controllers

The number of devices that VSC for VMware vSphere interprets as controllers but with which VSC for VMware vSphere is unable to

communicate. This could include controllers that are powered down, that do not have an IP address that is reachable from the VSC for VMware vSphere server, or that are not IBM N series storage controllers.

Number of Skipped Controllers

The number of controllers for which a user chose to skip credentials. Discovery is not attempted until the credentials are updated and the **Skipped** check box is cleared.

Number of Unmanaged Controllers

The number of controllers that are discovered but running a version of Data ONTAP software that VSC for VMware vSphere cannot manage.

Actions

The following actions are available on the Discovery Status panel:

Set Default Controller Credentials

Sets the default storage controller credentials. When a controller is discovered, VSC for VMware vSphere tries to connect using the default credentials. You can enter the credentials using the Modify Credentials dialog box (see the *Overview panel fields and commands* section for more information).

If you have controllers that have already been discovered but have a status of **Authentication Failure** or **SSL is not configured**, run an Update after setting new default credentials to re-try discovery using the new default credentials.

Note: The Monitoring and Host Configuration capability manages credentials for the Optimization and Migration capability and Provisioning and Cloning capability. You cannot enter separate credentials for those two capabilities.

Provisioning and cloning datastores and virtual machines

The Provisioning and Cloning capability of VSC for VMware vSphere enables you to provision datastores and quickly create multiple clones of virtual machines in the VMware environment.

The tasks you can perform with the Provisioning and Cloning capability include the following:

- Create clones of virtual machines and place them in new or existing datastores
- Create, resize, or delete datastores
- Apply guest customization specifications and power up new virtual machines
- Run deduplication operations
- Monitor storage savings
- Reclaim space on virtual machines stored in NFS datastores
- Redeploy virtual machines from a baseline image
- Replicate NFS datastores across sites
- Import virtual machines into virtual desktop infrastructure connection brokers and management tools

Tips for working with Provisioning and Cloning

There are several methods for accessing the Provisioning and Cloning features.

To manage datastores and clone virtual machines, right-click an object in the Inventory panel of the vSphere Client and select **IBM N series > Provisioning and Cloning**. Next you must right-click the correct object for the task you want to perform:

- To create clones, right-click a virtual machine or template.
- To provision datastores, right-click a datacenter, cluster, or host.

To manage controllers and connection brokers, replicate datastores, or redeploy clones, click the **Inventory** button in the navigation bar, and then select **Solutions and Applications > IBM N series**.

- To add, remove, or modify properties of storage controllers, select **Storage controllers**.
- To add or remove connection broker definitions, select **Connection brokers**.
- To clone NFS datastore templates to multiple target sites, select **DS Remote Replication**.
- To redeploy virtual machines, select **Redeploy**.

Cloning and managing virtual machines

You can use the Provisioning and Cloning capability to clone virtual machines, manage connection brokers, redeploy clones locally, and reclaim unused space on virtual machines.

Cloning virtual machines

You can use the Provisioning and Cloning capability to create theoretically thousands of virtual machine clones and hundreds of datastores at one time. In practice, however, multiple executions of fewer requests are recommended. The ideal size of the requests depends on the size of the vSphere deployment and the hardware configuration of the vSphere Client managing the ESX hosts.

Before you begin

Before you perform a cloning operation, it is a good practice to enable the NFS Plug-in for VMware VAAI. This plug-in is available from the N series support website (accessed and navigated as described in Websites). After you get the plug-in and place it in the correct directory, you can install it using the Monitoring and Host Configuration capability.

The following restrictions apply to this feature:

- The cloned virtual machine always has one virtual CPU (vCPU) no matter how many vCPUs the source virtual machine has.
- If you attempt to clone a virtual machine that has been functionally aligned using the Optimization and Migration capability, the clone will be misaligned. The Provisioning and Cloning capability warns you when you attempt to clone a functionally aligned virtual machine. This is because a functional alignment uses a prefix to get the virtual machine to align on the correct boundary. As a result, the virtual machine performs as though it has been aligned, but no changes have been made to the hard disk to ensure that the virtual machine is aligned to the storage system.
- You cannot use the cloning feature when the target virtual machine is being used by either the Backup and Recovery capability or the Optimization and Migration capability.

Procedure

1. In the vSphere Client Inventory, right-click a virtual machine or template and select **IBM > Provisioning and Cloning > Create rapid clones**. Cloning completes faster if the virtual machine or template is powered down. The Create Rapid Clones Wizard launches.
2. In the Storage Controller details window, select the target storage controller for the new clones from the drop-down list and also specify:
 - If you are using Data ONTAP operating in Cluster-Mode, specify a Vserver. The Provisioning and Cloning capability provides a drop-down list of the available Vservers.
 - If you are using Data ONTAP operating in 7-Mode, specify a vFiler. To identify vFiler units for Provisioning and Cloning operations, select the **Set vFiler Context** check box and select a unit from the drop-down list.
3. If the source virtual machine or template has snapshot copies, the Clone Source window displays. Select the template as the source for the new clones.

Note: If you specify a snapshot as the source, the Provisioning and Cloning capability clones the configuration file settings for the virtual machine, not the virtual machine itself.
4. In the Clone destination window, select the destination for the new clones. If you want to specify the virtual machine folder for the new clones, select the check box at the bottom of the window.

Note: If the destination is a datacenter, cluster, vApp, or resource pool, new clones are distributed across available servers as quickly as possible. New clones might be unevenly distributed, with more clones created on faster servers.

5. If selected, identify the virtual machine folder in which to place clones in the Virtual machine folder window.
6. In the Disk format window, select the disk format for the new clones.
7. In the Virtual machine details window, provide details for each virtual machine.

- To specify details manually, select the **Specify VM details** radio button.
- To import details from a file, select the **Import VM Details** radio button.

You can import the following virtual machine details from a .csv file:

- Non-contiguous virtual machine names
- Guest customization specifications
- virtual machine name as computer name (if guest customization specification is provided)
- Power-on setting

The file must contain the following fields: *cloneName*, *customSpecName*, *useVmNameAsPcName*, *powerOn*.

Fields 2 (*customSpecName*) and 3 (*useVmNameAsPcName*) are optional. Blank lines and lines that start with a hash (#) are ignored. The following is an example of a virtual machine details file:

```
# This is a sample VM details file
```

```
vm1, dewey, true, false
vm2, , , true
vm3, customSpecA, false, true
```

- **Create new datastores?** Creates new datastores for the virtual machine clones. Enabled for users with role of create or higher.
- **Import into connection broker?** Automatically imports clone data into a VMware View Server, or creates a .csv file for Citrix XenDesktop manual import. The .csv file is created in the directory [VSC_home]\etc\kamino\exports\xenDesktop_timestamp.csv, where [VSC_home] is your VSC for VMware vSphere installation directory.

Note: XenDesktop 5 does not support importing a virtual machine from vApps.

- **Connection broker:** Select the desired output type.
- **Virtual Processors:** Select number of virtual processors to apply to the new virtual machines.
- **Memory Size (MB):** Enter the amount of memory to apply to new virtual machines.
- **Upgrade hardware version?:** If clone source was created on ESX 3.5 or 4.x host and the destination is on a newer version of ESX, allows upgrade of clones to the new hardware version.
- **Clone operation details path:** (Import VM details only) Enter a file name or browse to the .csv file containing clone names.
- **Number of clones:** (Specify VM details only) For new datastores, maximum 250. For existing, maximum depends on available space. Number of clones must be evenly divisible by number of datastores being created.

Note: Success of 2,000 or more virtual machines depends on the size and performance of the vCenter Server.

- **Clone name:** (Specify VM details only) Prefix for each clone. By default, the clone number is placed at the end of the clone name. To force clone number to a different position, use `%CLONE_NUM%` where you want the number to appear. For example, `new%CLONE_NUM%clone`.
- **Starting clone number:** (Specify VM details only) -Maximum of eight digits.
- **Clone number increment:** (Specify VM details only) Increment clone numbers by 1, 2, 3, 4 or 5.
- **Power on?** (Specify VM details only) Select the check box if you want all the virtual machines to power on when the operation completes.

If you do not select this option, the Provisioning and Cloning capability leaves the clones powered off after the cloning process completes.

- **Stagger VM booting** (Specify VM details only) This option is available only when you select the **Power on?** option.

When you select this option, the Provisioning and Cloning capability staggers the start-up of the clones. In the **VMs per minute** box, you must supply an integer value indicating the number of clones per minute that you want the capability to power on.

Depending on your system setup and the number of clones you created, it is a good practice to stagger starting the clones so that you do not overwhelm your system. Having a large number of virtual machines start at once can slow down your system. The value you supply for this feature depends on your system environment at the time you perform the clone operation and how many clones you are creating.

Note: The Provisioning and Cloning capability keeps track of the total number of clones that should have powered on based on the number of minutes that have elapsed since the cloning operation completed. If a problem prevents the Provisioning and Cloning capability from starting the clones on schedule, it uses this cumulative total to specify how many clones to power on once the delay ends. Depending on the length of the delay and the number of clones per minute that you specified, any delay could result in the Provisioning and Cloning capability powering on a large number of clones at one time. After that, though, it only starts the specified number of clones per minute.

- **Apply customization specification?** (Specify VM details only) Applies a pre-defined specification to the new virtual machines. Select from specifications used for the native cloning process.
8. If you requested a new datastore, click one of the blue links in the Datastore creation window: **Create NFS datastore(s)** or **Create VMFS datastore(s)**. To continue without creating a new datastore, click **Next**.

You can create both NFS and VMFS datastores for the new virtual machines. Any restrictions for your configuration appear at the bottom of the window.

Note: Maximum VMFS datastore size and maximum number of NFS datastores depend on your version of VMware vSphere. In mixed version environments, maximums revert to earlier limits.

9. In the pop-up window, specify details for the new datastores.
- **Protocol:** (VMFS only) FCP or iSCSI.

- **Number of datastores:** Maximum 256. Number of clones must be evenly divisible by number of datastores.
- **Datastore name:** (single datastores only) Use the default or replace with a custom name.

Note: For multiple datastores, the golden volume name (for NFS) or base name (for VMFS) is used here and in the Summary window as the datastore name.

- **Size:** Maximum depends on the controller and space available. For details, see the *Data ONTAP Storage Management Guide* for your Data ONTAP release.
- **Create new volume container:** (VMFS only) Create a volume with the same name as the LUN. If a volume with that name already exists, the volume name is appended with a number; for example, **Volname01**.
- **Volume:** (VMFS only) Select available volume from drop-down list.
- **Aggregate:** Select available aggregate from drop-down list.
- **Thin provision:** Sets space reserve to none and disables space checks.

Note: Cloning and datastore creation can fail if the size request uses too much of the aggregate. Capacity is not reserved for individual datastores. Instead, the aggregate is treated as a shared pool with capacity used as each datastore requires it. By eliminating unused but provisioned storage, more space is presented than is available. It is expected that the datastores will not utilize all provisioned storage at once.

- **Block size:** (VMFS only) Select block sizes. For VMFS-5, block size is fixed at 1 MB.
 - **Auto-grow:** (NFS only) When space is needed, automatically expands the datastore by increment you specify, up to size limit you specify.
 - **Grow increment:** (NFS only) Amount of storage added to datastore each time space is needed.
 - **Maximum datastore size:** (NFS only) Limit at which Auto-grow stops.
 - **Datastore cluster:** (vCenter 5 or later) Select a datastore cluster in which to add the datastore. Requires SDRS on the vCenter Server.
 - **Set datastore names?** (multiple datastores only) Enables modification of default datastore names.
 - **Group name:** (VMFS) or Golden volume name (NFS) Use defaults or replace with custom names.
10. In the Datastore selection window, select the datastore to house the new virtual machines. You can clone the VMDK files that comprise the virtual machine to different datastores.
 - To place all virtual machines in a single datastore, select the datastore and click **Next**.
 - To distribute virtual machine files across multiple datastores, click **Advanced**. In the new window, select a virtual machine file, open the corresponding Datastore pull-down list, and select the datastore to house that file. Repeat for each virtual machine file, then click **Next**
 11. If you specified a connection broker format, enter the details for the connection broker operation.

Option	Description
VMware View Server	<p>For VMware View Server, clone data is imported into View Server at the end of the clone operation. You can specify unique desktop names and select from existing desktop pools.</p> <p>Note: Refer to the <i>View Manager Administration Guide</i> at http://www.vmware.com/pdf/view40_admin_guide.pdf for details on desktop types and access modes.</p>
Citrix XenDesktop	<p>A .csv file is created in the directory [VSC_home]\etc\kamino\exports, where [VSC_home] is your VSC for VMware vSphere installation directory, with the following format:</p> <p>XenDesktop 4</p> <p>[ADComputerAccount],[AssignedUser],</p> <p>[VirtualMachine],[HostID]</p> <p>XenDesktop 5</p> <p>[VirtualMachinePath],[ADComputerAccount],</p> <p>[AssignedUsers]</p> <p>See Importing virtual machines into XenDesktop for details on importing this file into Citrix XenDesktop.</p>

Note: XenDesktop 5 does not support importing from vApps.

- **View Server hostname or IP Address:** (VMware View Server only) Enter the hostname or IP address of the VMware View Server.
- **Connection name:** (Citrix XenDesktop 5.0 only) Enter the name of the VMware connection.
- **Domain name:** Enter a fully qualified domain name.
- **Domain username:** (VMware View Server only) Enter the domain user name.
- **Domain password:** (VMware View Server only) Enter the domain password.
- **New desktop pool:** Allows you to import new virtual machine clones into multiple new View Server pools.
 - **Number of pools:** Enter the number of pools to create.
 - **Name pools automatically?** Select the checkbox to use the specified pool name as prefix for each pool name.
 - **Distribute VMs evenly?** Select the checkbox to distribute new virtual machines evenly across new desktop pools.

To distribute varying quantities of virtual machines to pools, do not select the **Distribute VMs evenly** checkbox. Instead, enter the desired values directly in the table. The sum of the new distribution must equal the number of new clones.

To provide unique names to the new pools, do not select the **Name pools automatically** checkbox. Instead, enter desired names directly in the table.

- **Desktop pool type:** Select access type for new desktop pools:
 - **Dedicated:** (View Server 4.5 and higher) gives users connection to the same desktop for every session.
 - **Floating:** (View Server 4.5 and higher) allocates desktops dynamically.

- **Existing desktop pool:** (View Server only) Allows you to select from existing pools on the specified View Server.

If you have changed the View server or updated credentials, click Refresh to update the list of pool names.

- **Desktop pool name:** Select an existing pool on the specified View Server from the drop-down list.

12. Review the summary page and click **Apply** to proceed. To return to previous pages and modify settings, click **Back**.

Results

The Recent Tasks pane of vSphere Client is populated as clone creation proceeds.

A list of new virtual machines is written to a .csv file in the directory [VSC_home]\etc\kamino\exports, where [VSC_home] is your VSC for VMware vSphere installation directory. The filename is in the following format: import_generic_date_time.csv

Managing connection brokers

You can use the Connection brokers panel to view and manage the connection brokers available for importing clone data at the end of the clone operation.

- For VMware View Server, clone data is imported into View Server at the end of the clone operation.
- For Citrix XenDesktop, a .csv file is created in the directory c:\program files\ibm\virtual storage console\etc\kamino\exports. See Importing virtual machines into XenDesktop for details.

Adding connection brokers

You can add connection brokers to the Provisioning and Cloning capability.

Procedure

1. Select the **IBM System Storage N series** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **Connection brokers**.
3. Click the blue **Add Connection Broker** link.
4. In the Add Connection Broker window:
 - a. **Connection broker** - Select the name and version of the desired connection broker.
 - b. **Domain** - Enter the domain containing the connection broker.
 - c. **Connection name** (XenDesktop 5.0 only) - Enter the name given the Citrix XenDesktop 5.0 connection.
 - d. **Hostname or IP Address** (VMware View Server only) - Enter the connection broker hostname or IP address.
 - e. **Username** (VMware View Server only) - Enter the domain user name.
 - f. **Password** (VMware View Server only) - Enter the domain password.

Removing connection brokers

You can remove a connection broker from the list of available brokers.

Procedure

1. Select the **IBM System Storage N series** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **Connection brokers**.
3. Click the blue **Remove Connection** broker link.
4. Click **Yes** to confirm.

Redeploying clones (locally)

You can reset to their original state all virtual machines that are based on a selected gold virtual machine or template. You can also propagate changes made in the original gold virtual machine or template to all of its clones and, optionally, reapply customization specifications as well.

Before you begin

This feature has the following requirements:

- It is available only for virtual machines that reside entirely on NFS destinations.
- It only applies templates to local virtual machines.
To replicate templates from a vCenter to subordinate vCenters across multiple sites, see Replicating remote datastores.
- You cannot use this feature when the target virtual machine is in use by the Backup and Recovery capability or the Optimization and Migration capability.

Procedure

1. Select the **IBM System Storage N series** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **Redeploy**.
3. Select the virtual machine or template to use, then click the **Redeploy...** link on the right side of the window.

Note: To refresh the list, click the **Update table** link.

4. In the Select clones window, select the check boxes for the clones to redeploy.
5. In the Redeploy clones window, specify the settings to apply to the redeployed clones.

Note: If the selected virtual machines are individual desktops or a dedicated pool in VMware View 4.5 and higher, the redeploy may result in loss of user data.

- a. **Power on?:** Powers on all the virtual machine after the operation completes.
If you do not select this option, the Provisioning and Cloning capability leaves the virtual machines powered off after the cloning process completes.
- b. **Stagger VM booting:** This option is available only when you select the **Power on?** option.

If you select this option, the Provisioning and Cloning capability staggers the start-up of the cloned virtual machines. In the **VMs per minute** box, you must supply an integer that indicates how many virtual machines the Provisioning and Cloning capability should start each minute.

Depending on your system setup and the number of clones you created, it is a good practice to stagger starting the virtual machines so that you do

not overwhelm your system. Having a large number of virtual machines start at once can slow down your system. The values you supply for this feature depend on your system environment at the time you perform the clone operation and how many clones you are creating.

Note: The Provisioning and Cloning capability keeps track of the total number of virtual machines that should have powered on based on the number of minutes that have elapsed since the cloning operation completed. If a problem prevents the Provisioning and Cloning capability from starting the virtual machines on schedule, it uses this cumulative total to specify how many virtual machines to power on when the delay ends. Depending on the length of the delay and the number of clones per minute that you specified, any delay could result in the Provisioning and Cloning capability powering on a large number of virtual machines at one time. After that, though, it only starts the specified number of virtual machines per minute.

- c. **Apply customization specification?:** From the drop-down list of available customized specifications, select the one that you want to apply to these clones.
- d. **Use the virtual machine name as the computer name?** If you are using a customization specification with a custom sysprep answer file, select this check box to insert the virtual machine name in the answer file.
- e. Review the summary of choices for this operation and click **Apply** to proceed.

Removing a baseline

The Redeploy window displays available baselines that have been used to create clones. You can remove baselines from this list.

Procedure

1. Select the **IBM System Storage N series** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **Redeploy**.
3. Select the baseline and click the blue **Remove** link on the right side of the window.

Reclaiming space on virtual machines

You can use the Reclaim space feature to find free clusters on NTFS partitions and make them available to the operating system.

Before you begin

The Reclaim space feature allows Data ONTAP to use space freed when data is deleted in guest operating systems.

This feature has the following requirements:

- VMDKs attached to the virtual machine must be on NFS-backed datastores.

Note: The Reclaim space feature is not supported if the NFS datastore is backed by a qtree on a vFiler unit.

- VMDKs must have NTFS partitions.

Note: If the VMDK is unpartitioned or FAT, the Provisioning and Cloning capability incorrectly lists the disk as having an NTFS partition after the task

completes and displays a "Yes" in the "Has NTFS partition(s)?" column. Even though the VMDK now appears to be partitioned, it is still unpartitioned or FAT, and you cannot reclaim space on it.

- ISOs mounted to the virtual machine must be contained in an NFS datastore.
- Storage systems must be running Data ONTAP 7.3.4 or later.
- You should have the VMware guest tools installed.
- When the Reclaim space feature is running, you must not power on the virtual machine.
- You cannot use the cloning feature when the target virtual machine is being used by either the Backup and Recovery capability or the Optimization and Migration capability.

Procedure

1. Right-click a datastore or virtual machine and select **IBM > Provisioning and Cloning > Reclaim space**.
2. Click **OK**. If the virtual machine is powered on, the Reclaim space feature powers it off. After the process completes, the Reclaim space feature returns the virtual machine to its previous state.

Note: If you are using this feature when the virtual machine is powered on, make sure you have the guest operating system tools installed. Without these tools, the Reclaim space feature does not work when it has to power down the virtual machine

If you do not want to install these tools, then you should power down the virtual machine before running the Reclaim space feature.

Importing virtual machines into XenDesktop

You can manually import virtual machines into XenDesktop.

To create a Citrix XenDesktop import file, select one of the Citrix XenDesktop versions as the connection broker in the Virtual Machine Details window of the Create Rapid Clones Wizard.

The Provisioning and Cloning capability creates the following XenDesktop import file, where [VSC_home] is your VSC for VMware vSphere installation directory:

```
[VSC_home]\etc\kamino\exports\xenDesktop_timestamp.csv
```

The file contents are in the following format:

```
[ADComputerAccount],[AssignedUser],[VirtualMachine],[HostID]
```

Note: XenDesktop 5 does not support importing from vApps.

Importing the file into XenDesktop 4

You can manually import a Citrix XenDesktop import file into XenDesktop 4.

Procedure

1. Copy the import file to the XenDesktop system.
2. Using the Citrix Access Management Console, choose to create a new desktop group. A wizard launches.

3. Follow the wizard prompts to the fifth panel (Virtual Desktops page), and select the option to import the desktops from a .csv file.
4. Browse to the import file and click **OK**.

Importing the file into XenDesktop 5

You can manually import a Citrix XenDesktop import file into XenDesktop 5.

Procedure

1. Copy the export file to the XenDesktop system.
2. In Desktop Studio, create a new catalog or modify an existing one in Desktop Studio. A wizard launches.
3. Follow the wizard prompts and select the option to import existing virtual machines.
4. Browse to the import file and click **OK**.

Managing storage controllers

You can view usage and deduplication statistics, and manage volume settings and network interfaces with the Provisioning and Cloning capability.

(Data ONTAP operating in 7-Mode) Viewing storage controller details

You can view usage and deduplication statistics for a storage controller's aggregates, volumes, and LUNs.

About this task

These steps apply to storage systems that are running Data ONTAP operating in 7-Mode.

Procedure

1. Select the **IBM System Storage N series** icon for the vCenter Server.
2. In the left pane under Provisioning and Cloning, click **Storage controllers**.
3. Right-click a storage controller and select **View Storage Details**.

(Data ONTAP operating in 7-Mode) Removing or adding network interfaces, volumes, and aggregates

Volumes and aggregates created outside the Provisioning and Cloning capability must be added in the Resources window to be available for provisioning and cloning operations. You can also use this window to restrict or expand the available network interfaces, volumes and aggregates. This feature is available only on storage systems running Data ONTAP operating in 7-Mode.

About this task

Note: If you are running Data ONTAP operating in Cluster-Mode, the recommended best practice is to create a Vserver that is suited to the needs of the users. That way you use the security and RBAC capabilities provided by Data ONTAP to control what is available to VSC for VMware vSphere and the ESX/ESXi hosts.

Procedure

1. Select the **IBM System Storage N series** icon for the vCenter Server.
2. In the left pane under Provisioning and Cloning, click **Storage controllers**.
3. Right-click a storage controller and select **Resources**. Or click the blue **Resources...** link. All available network interfaces, volumes, and aggregates are presented.
 - a. To restrict or expand availability, select components and click the directional buttons to move them to the desired column.
 - Items in the left column will *not* be used.
 - Items in the right column will be used.

Note: New volumes and aggregates created outside the Provisioning and Cloning capability appear in the left column as unused.

 - b. To lock these settings and require storage system credentials for modifications, select the **Prevent further changes** check box at the bottom of the window.
4. Click **Save** when finished.

Managing volume settings

You can establish advanced settings for new volumes on the storage controller.

Procedure

1. Select the **IBM System Storage N series** icon for the vCenter Server.
2. In the left pane under Provisioning and Cloning, click **Storage Controllers**.
3. Right-click a storage controller and select **Settings**. Or click the blue **Settings...** link. If storage controller properties have been locked to prevent changes, you will be prompted for the username and password for that storage controller. You must enter them before you can change volume settings.
4. Set the following options for new volumes created on this storage controller.

If a thin provisioned LUN is deployed into a FlexVol with volume autogrow or snapshot autodelete disabled, it is possible to overcommit the LUN to the volume. This creates an out of space condition.

Use these advanced options to modify the FlexVol efficiency settings to match those on the LUN being deployed:

- **Create a new volume for a new LUN** - create a FlexVol with the same name as the LUN. If a volume with that name already exists, the volume name is appended with a number; for example, Volname01.

Note: If you want the volume container to be resized with the LUN, select this option.

- **Reserve space for volumes that contain thin provisioned LUNs** - results in a thin LUN in a thick volume when a thin LUN is chosen.
 - **Thin provision volume clones** - sets space reservation policy to thin provisioning for clones created from this volume.
 - **Delete a volume if the last LUN in it has been deleted** - destroy volume when its last LUN is deleted.
 - **Buffer space between volume and LUN (GB)** - Amount of additional capacity in a volume that contains a LUN based datastore.
5. Click **Save** when finished.

Managing datastores

You can use the Provisioning and Cloning capability to replicate datastores to remote sites, provision, mount, resize, and destroy datastores, and manage deduplication on datastores.

(Data ONTAP operating in 7-Mode) Replicating datastores to remote sites

The Datastore Remote Replication feature uses Asynchronous SnapMirror to clone NFS datastores from a source vCenter to one or more remote vCenter sites. This feature is available only on storage systems running Data ONTAP operating in 7-Mode.

Setting up a replication target establishes the SnapMirror schedule and creates a FlexVol volume at the target site that is used as the SnapMirror destination. After the initial SnapMirror transfer is complete, datastore synchronization takes place at the target site. Subsequently, the SnapMirror schedule continues to push incremental updates to the target site so the data is available for any future datastore synchronization operations. After the initial setup, a datastore is not set up at the target site until a manual synchronization is performed from the Datastore Remote Replication panel.

Note: The target vCenter version must be the same as or later than the source vCenter version. Otherwise, synchronization succeeds but virtual machines registered on the target are invalid.

The Datastore Remote Replication panel displays the following fields:

- **Sources** - Datastores available to be replicated. Available datastores reside on storage controllers that have SnapMirror licenses.

Note: vFiler context-based NFS datastores are not supported.

- **vCenter** - vCenter containing the hosts where the datastore's virtual machines will be registered.
- **Storage Controller** - Storage controller where source datastore will be replicated. Available storage controllers meet the following requirements:
 - have a valid SnapMirror license
 - are not vFiler units
 - have the same Data ONTAP version as the storage controller where the source datastore resides (SnapMirror requirement)
- **Aggregate** - Aggregate where the source datastore will be replicated.
- **Datastore** - Name for the target datastore. The name of the FlexVol volume created as the SnapMirror destination will have `_distribution` appended to it. During synchronization, a clone of the FlexVol volume is created using the name specified in this field.
- **SnapMirror Lag** - Difference between the current time and the timestamp of the snapshot copy last transferred to the destination successfully.
- **SnapMirror Schedule** - Minute, hour, month, and day defined for Asynchronous SnapMirror update. Incremental data is sent from source datastore to a FlexVol volume at the target destination and held until the next Synchronize is performed.

Note: This feature replicates datastores from a vCenter to subordinate vCenters across sites. To apply templates to local virtual machines, refer to Redeploying clones (locally).

Setup Replication

You can define a new datastore replication relationship using the Provisioning and Cloning capability.

About this task

Note: After establishing a remote replication relationship, if you uninstall and re-install the destination vCenter, you will have to recreate the relationship.

Procedure

1. Select the **IBM System Storage N series** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **DS Remote Replication**.
3. Select the blue **Add** link.
4. In the Source Datastore window, select a source datastore. Qualified datastores reside on **IBM** storage controllers with SnapMirror licenses.
5. In the Target vCenter window, add or select the vCenter containing the target datastore. You can also use this window to remove a selected vCenter from the list.
 - Click **Add vCenter** if the list is empty or to add another vCenter for use as a target. Enter the **vCenter name**, **Username**, and **Password**.
 - Click **Remove** to remove a selected vCenter from the list.
6. In the Target Destination window, identify the destination where the source datastore will be distributed.
 - a. First select the host, cluster, or datacenter in which to register the new virtual machines.
 - b. In the Target Storage Controller list box, select the storage controller where the target datastore will be created.
 - c. In the Aggregate list box, select the aggregate where the target datastore will be created. Available aggregates contain space equivalent to or greater than the source volume, and have the same bit type as the source.
 - d. In the **Target Datastore** box, enter a name for the datastore to contain the new virtual machines.
7. In the Network Mapping window, establish the network mappings between source and destination virtual machines. Select a destination network from the drop-down for each source network.
8. In the SnapMirror Schedule window, establish the schedule for Asynchronous SnapMirror update.

Incremental data is sent from the source datastore to a FlexVol volume at the target destination and held until the next Synchronize is performed.

Note: All fields accept comma-separated sequences as well as ranges. For example, 1, 2, 3, 4-7, 8, 10. To represent all valid digits, enter an asterisk (*).

- a. **Minute** - Enter the minute past the hour for Asynchronous SnapMirror update to begin. Accepts 0 - 59.
- b. **Hour** - Enter the hour of the day for Asynchronous SnapMirror update to begin.

Example

For example, enter 3 for 3 a.m., and 15 for 3 p.m.

Accepts 0 - 23.

- c. **Month** - Enter the month of the year for Asynchronous SnapMirror update to occur. Accepts 1 - 12.
 - d. **Day of Week** - Enter the day of the week for Asynchronous SnapMirror update to occur. Accepts 0 - 6.
 - e. **Day of Month** - Enter the day of month for Asynchronous SnapMirror update to occur. Accepts 1 - 31.
9. Review the summary and click **OK** to proceed or **Cancel** to exit the wizard without setting up a replication relationship.

Remove a datastore replication relationship

This feature places the SnapMirror relationship between source and target in Broken-off status. SnapMirror activity ends, but the volume remains at the target site. The deleted relationship remains in the output of the CLI command **snapmirror status** until you delete all snapshots created on the source volume as part of the Datastore Remote Replication process.

Procedure

1. Select the **IBM System Storage N series** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **DS Remote Replication**.
3. Select a target and then select the blue **Remove** link.

Synchronize

You can manually initiate datastore replication to the designated targets, or write stored incremental updates to the target.

Before you begin

This action powers off all virtual machines in the source datastore.

Procedure

1. Select the **IBM System Storage N series** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **DS Remote Replication**.
3. To replicate templates or template changes, right-click a target and select **Synchronize**.

Define or modify a SnapMirror schedule

You can modify the SnapMirror schedule using the Provisioning and Cloning capability. The SnapMirror schedule defines the schedule for Asynchronous SnapMirror updates. Incremental data is sent from the source datastore to a FlexVol volume at the target destination and held until the next Synchronize is performed.

About this task

Note: All fields accept comma-separated sequences as well as ranges. For example, 1, 2, 3, 4-7, 8, 10. To represent all valid digits, enter an asterisk (*).

Procedure

1. Select the **IBM System Storage N series** icon for the vCenter Server.

2. In the navigation pane under Provisioning and Cloning, click **DS Remote Replication**.
3. Click the blue **SnapMirror Schedule** link.
4. In the **SnapMirror Schedule** window:
 - a. **Minute** - Enter the minute past the hour for Asynchronous SnapMirror update to begin. Accepts 0 - 59.
 - b. **Hour** - Enter the hour of the day for Asynchronous SnapMirror update to begin. For example, enter 3 for 3 a.m., and 15 for 3 p.m. Accepts 0 - 23.
 - c. **Month** - Enter the month of the year for Asynchronous SnapMirror update to occur. Accepts 1 - 12.
 - d. **Day of Week** - Enter the day of the week for Asynchronous SnapMirror update to occur. Accepts 0 - 6.
 - e. **Day of Month** - Enter the day of month for Asynchronous SnapMirror update to occur. Accepts 1 - 31.

Refreshing the display

You can refresh the Datastore Remote Replication display to show the most recent changes.

Procedure

1. Select the **IBM System Storage N series** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **DS Remote Replication**.
3. Click the blue **Refresh** link.

Provisioning datastores

You can create new datastores at the datacenter, cluster, or host level. The new datastores appear on every host in the datacenter or the cluster. For Data ONTAP operating in Cluster-Mode, you can either select the cluster and choose the Vserver through which to provision the new volume, or provision directly to the Vserver.

About this task

You cannot provision NFS datastores to a direct-connect vFiler unit. A direct-connect vFiler unit cannot add volumes to itself. To enable NFS provisioning, you must add vfiler0 and then select the vFiler unit to which you want to provision storage.

There are also limitations with direct-connect Vservers:

- To provision datastores on direct-connect Vservers, you must create a new role and a new user with the required privileges to provision. The default vsadmin role assigned to Vservers does not contain the volume efficiency commands needed by the Provisioning and Cloning capability.
- Server privileges are not visible to the Provisioning and Cloning capability. When connecting directly to a Vserver, Provisioning and Cloning operations might begin but fail later due to insufficient privileges.
- Datastore creation is not supported on striped aggregates. When you add a direct-connect Vserver, any striped aggregates associated with that vServer appear as available. Provisioning on those striped aggregates will fail.

Procedure

1. In vSphere Client Inventory, right-click a datacenter, cluster, or host and select **IBM > Provisioning and Cloning > Provision datastores**.
2. In the Storage Controller details window:
 - a. Select the target physical storage controller, vFiler unit, cluster, or Vserver.
 - b. If you selected a cluster as the target, select the **Vserver** through which to provision datastores onto the cluster.
 - c. To identify vFiler units for provisioning and cloning operations, select the **Set vFiler Context** check box and select a unit from the **vFiler** drop-down list.
3. In the Datastore type window, select the datastore type.
4. In the Datastore details window, specify the following details for the new datastore.
 - **Protocol:** (VMFS only) **FCP** or **iSCSI**.
 - **Size:** Maximum datastore size depends on the controller and space available. For details, see the *Data ONTAP Storage Management Guide* for your Data ONTAP release.
 - **Datastore name:** Use the default or replace with a custom name.
 - **Create new volume container:** (VMFS only) Create a FlexVol to contain the new datastore.

Note: If you want the volume container to be resized with the LUN, select this option.

- **Volume:** (VMFS only) Select an available volume from drop-down list.
- **Aggregate:** Select an available aggregate from the drop-down list.
- **Thin provision:** Sets space reserve to none, and disables space checks.

Important: Cloning and datastore creation can fail if your size request uses too much of the aggregate. Capacity is not reserved for an individual datastore. The aggregate is treated as a shared resource pool, where capacity is consumed as each datastore requires it. By eliminating unused but provisioned areas of storage, more space is presented than is actually available. It is expected that all datastores will not use all of their provisioned storage at the same time.

- **Block size:** (VMFS only) Select an available block size from the drop-down list. For VMFS-5, block size is fixed at 1 MB.
 - **Auto-grow:** (NFS only) - If more space is required, automatically expands the datastore by the increment you specify, up to the size limit you specify.
 - **Grow increment:** (NFS only) Amount of storage added to datastore each time space is needed.
 - **Maximum datastore size:** (NFS only) Limit at which Auto-grow stops.
 - **Datastore cluster:** (vCenter 5 only) Select a datastore cluster in which to add the datastore. Requires SDRS on the vCenter Server.
5. Review the datastore configuration summary and click **Apply** to begin provisioning the new datastore.

Mounting datastores

You can add currently mounted datastores to a new host.

About this task

Note: This feature is unavailable when the target virtual machine is in use by the Backup and Recovery capability or the Optimization and Migration capability.

Procedure

1. In the vSphere Client Inventory, right-click the new host and select **IBM > Provisioning and Cloning > Mount datastores**.
2. Select the datastores to mount on the host. The drop-down list contains datastores on the cluster containing the new host. Use CTRL-click to select multiple datastores.
3. Click **OK**.

Managing deduplication

You can enable deduplication on a selected datastore to optimize space utilization.

Before you begin

For more about deduplication, refer to your *Data ONTAP Data Protection Online Backup and Recovery Guide*.

Procedure

1. In the vSphere Client Inventory, right-click a datastore and select **IBM > Provisioning and Cloning > Deduplication management**.
2. Select the appropriate check boxes for your configuration. The Deduplicated column presents previously used space saved by this feature.
 - **Enable/Disable** toggles the deduplication feature on or off.
 - **Start** begins deduplication from the last marker position.
 - **Scan** begins deduplication at the beginning of the volume.

Resizing datastores

You can increase or decrease NFS datastore sizes. VMFS datastore sizes can be increased but not decreased.

About this task

If you want the volume container to be resized with the LUN, go to **Storage Controllers > Settings**. Select the **Advanced** checkbox, and select the option to **Create a new volume for a new LUN**. Otherwise, the container volume will not be resized with the LUN.

Note: This feature is unavailable when the target virtual machine is in use by the Backup and Recovery capability or the Optimization and Migration capability.

Procedure

1. In the vSphere Client Inventory, right-click a datastore and select **IBM > Provisioning and Cloning > Resize**.
2. Enter the new datastore size and click **OK**.
3. Click **Yes** to confirm the operation.

Destroying datastores

You can permanently destroy a datastore, including all virtual machines, datastores, Snapshot copies and FlexVolumes.

About this task

The destroy datastore feature performs the following actions:

- destroys all virtual machines in a datastore
- unregisters and detaches the datastore from the vSphere Client environment
- frees the space on the storage controller

Note: This feature is unavailable when the target virtual machine is in use by the Backup and Recovery capability or the Optimization and Migration capability.

Procedure

1. In the vSphere Client Inventory, right-click a datastore and select **IBM > Provisioning and Cloning > Destroy**.
2. Click **OK** to proceed.

About multiple datastores (NFS only)

When you destroy the last datastore of a golden volume, the golden volume is destroyed only if you select the option to **Delete a volume if the last LUN in it has been deleted** for that volume.

When you create more than one new NFS datastore, the Provisioning and Cloning capability first creates a golden volume on the controller. The datastores that are then attached to the vSphere Client are FlexClones of the golden volume. Because the FlexClones share the storage with the golden volume, this space is not wasted. When the last datastore (FlexClone) of a golden volume is destroyed, the golden volume is destroyed only if you select the option to **Delete a volume if the last LUN in it has been deleted** for that volume. For details, see Managing volume settings.

Provisioning and Cloning support files

There are several support files used and created by the Provisioning and Cloning capability.

Preferences File

The preferences file contains options that control Provisioning and Cloning capability operation.

This file is stored by default in the following location, where [VSC_home] is your VSC for VMware vSphere installation directory:

```
[VSC_home]\etc\kamino\kaminoprefs.xml
```

The preferences file contains the following options.

Type	Value	Description
log4j.config.file	log4j.properties	Filename of the log configuration file. This file contains parameters that determine the level of events that are written to the log file. The log configuration filename can be changed here in the kaminoprefs file, and the new configuration file must be stored in [VSC_home]\etc\kamino.
default.create.copyNvram File	false	Clones the VM BIOS file (.nvram). This option is necessary when using Citrix Provisioning server because the boot device order and PXE settings need to be cloned as well.
default.create.exportToAllVmKernelNics	false	Determines if the storage controller exports file will contain all IP addresses of all VMkernel ports. For example, if a dedicated storage network with IP address 192.168.0/24 and a second VMkernel network 10.10/16 on a public network for ISO mounting have this value set to true, the newly created datastores are exported to VMkernel addresses on both networks. The default behavior of false restricts the export to only the VMkernel addresses that share the same network as the IP address of the controller used to mount the datastore.
default.create.volume.option.fractional_reserve	0	Fractional reserve allows the system administrator to reduce the amount of guaranteed available reserve space. By doing this, the administrator can tune the amount of space guaranteed to be available for LUN overwrites based on application requirements and make better use of available volume space. It is adjustable between 0 and 100% of the size of the LUN.

Type	Value	Description
<code>default.create.volume.option.no_atime_update</code>	on	Disables updating of access time when files are read. Leave on for better performance in a heavy read traffic environment.
<code>default.create.volume.option.nosnap</code>	on	Disables Snapshot copies for new FlexClone volumes.
<code>default.create.volume.option.sis</code>	true	Enables deduplication on new volumes.
<code>default.create.volume.snapshotreserve</code>	0	Sets the snapshot reserve size for new volumes to the specified percentage.
<code>default.create.volume.lunSpaceBuffer</code>	5	Sets 5 GB as the amount of additional capacity in a volume that contains a LUN-based datastore.
<code>default.create.volume.useThickForVolThatContainLun</code>	false	Space reservation override for thin LUN-based datastore. Default volume space reservation for a thin LUN is thin. Setting this option to true results in a thin LUN in a thick volume when a thin LUN is chosen. Note: Applies only to Datastore Create. It has no effect on Datastore Clone.
<code>default.create.volume.useThinForVolClone</code>	true	Space reservation override for volumes created using FlexClone technology. Datastores will all be thin FlexClones unless this option is set to false. Note: Applies only to Datastore Clone. It has no effect on Datastore Create.

Type	Value	Description
default.destroy.checkPossibleCrossMount	true	If an administrator mounts a volume for use as a datastore to an ESX host using a different IP address or volume path than what was used to mount the volume to other ESX hosts, there is a potential for the Destroy feature to incorrectly report which VMs are present on the datastore. This is because the datastore UUID is generated by the ESX host based on the IP address and volume path. checkPossibleCrossMount addresses this issue and therefore should be left set to true. The only time this should be set to false is if the Destroy window stays in the "Loading..." state for an extended period of time and there is no chance of the issue described above.
default.destroy.destroyParentIfLastClone	true	When last FlexClone volume is destroyed, the golden volume is destroyed.
default.destroy.onlyOffline	false	By default, the Destroy feature takes offline and destroys the volume that was mounted as a datastore. Setting this option to true leaves the volume intact, just offline.
default.destroy.onlyUnreg	false	By default, the Destroy feature chooses the most efficient VM cleanup method. Setting this option to true forces it to unregister the VM in all cases.
default.destroy.destroyVolumeIfLastLun	false	When last LUN is removed, volume is not destroyed. Setting this option to true causes volume to be destroyed when last LUN is removed.
default.mount.volume.mountUsingHostname	Not present	If true, hostname is retrieved by reverse name lookup using system-configured lookup service. Hostname is then used to mount the datastore.

Type	Value	Description
default.restrict.nfs.mount.networks	None	<p>By default, any matching network between the controller and the ESX VMKernel is used to mount datastores. In some implementations, certain networks should not be used for mounting, even if they contain the required interfaces on both the controller and ESX hosts. This value prevents a network from being used to mount NFS datastores. The example below tells the Provisioning and Cloning capability to ignore the 192.168.2.0, 10.1.0.0, and 2001:5a0:400:200::0 networks when determining the IP address on the controller to use for the next NFS datastore:</p> <pre><entry key="default.restrict.nfs.mount.networks" value="192.168.2.0;10.1.0.0;2001:5a0:400:200::0"/></pre>
default.restrict.iscsi.mount.networks	None	<p>By default, any matching network between the controller and the ESX VMKernel is used to mount datastores. In some implementations, certain networks should not be used for mounting, even if they contain the required interfaces on both the controller and ESX hosts. This value prevents a network from being used to mount iSCSI datastores. The example below tells the Provisioning and Cloning capability to ignore the 192.168.2.0, 10.1.0.0, and 2001:5a0:400:200::0 networks when determining the IP address on the controller that will be used for the next iSCSI datastore.</p> <pre><entry key="default.restrict.iscsi.mount.networks" value="192.168.2.0;10.1.0.0;2001:5a0:400:200::0"/></pre>

Type	Value	Description
default.allow.nfs.mount.networks	all	Any matching network between controller and ESX VMKernel is used to mount NFS datastores. In some implementations, it is important to fence off the networks that should be used for NFS and those that should be used for iSCSI. This option allows only specified networks to mount NFS datastores. The network specified is that of the VMKernel interface, not an interface on the storage system. The all keyword means that all matching networks between the VMkernels and the storage system interfaces can be used for mounting NFS datastores. You can specify different networks here to enable mounting across different subnets.
default.allow.iscsi.mount.networks	all	Any matching network between controller and ESX VMKernel is used to map iSCSI datastores. In some implementations, it is important to fence off the networks that should be used for NFS and those that should be used for iSCSI. This option allows only the specified networks to map iSCSI-based VMFS datastores. The network specified is that of the VMKernel interface, not an interface on the storage system.
default.create.igroup.alua	yes	Configures the ALUA setting to apply to new iGroups.
default.option.copythreads	4	Controls number of simultaneous file copies. Maximum is 16.
default.option.powerOffThreads	10	Controls number of simultaneous VM power off operations. Maximum is 16.
default.option.checkAlignment	true	Controls whether alignment is checked on the source VM.

Type	Value	Description
default.option. fallBackToNdmpCopy	true	Specifies whether the clone operation should fall back to using NDMP copy if the flex-clone operation begins to copy blocks rather than clone them. This happens when a block has reached the max shared blocks limit.
default.option. fallBackThreshold	0	Specifies a threshold for copying blocks before falling back to using an NDMP copy. Blocks copied threshold is specified as a percentage. This is used if the default.option.fallBackToNdmpCopy preference is set to true. For example, setting this value to 20 will fall back to using NDMP copy if 20% of the blocks during the FlexClone operation have been copied. Default behavior is to fall back immediately when any blocks have been copied.
default.option. controllerValidationTimeout	120	By default, controller validation is cancelled if it takes longer than two minutes to validate a controller. In some large environments, it may be necessary to increase this timeout. The example below tells the Provisioning and Cloning capability to wait five minutes (300 seconds) before canceling the validation operation. <entry key="default.option.controllerValidationTimeout" value="300"/>

Provisioning and Cloning logs

The Provisioning and Cloning capability creates a log file that includes the capability version, build number, build date, and all event messages and errors.

The log file is created in the following location, where [VSC_home] is your VSC for VMware vSphere installation directory:

[VSC_home]\log\kamino.log

Each log file is rotated at a size of 8,192 KB. Up to nine log files are retained.

Log configuration file

The level of events recorded in the log file is determined by parameters set in the log configuration file.

The log configuration file is stored in the following location, where [VSC_home] is your VSC for VMware vSphere installation directory:

```
[VSC_home]\etc\kamino\log4j.properties
```

Modifying logging levels

You can change the level of events that are logged for the Provisioning and Cloning capability at runtime by modifying the log configuration file.

Edit the log configuration file [VSC_home]\etc\kamino\log4j.properties, where [VSC_home] is your VSC for VMware vSphere installation directory.

The new settings take effect immediately. It is not necessary to restart the VSC for VMware vSphere service.

Using a different log configuration file

You can define a different log configuration file by entering the new configuration file name in the kaminoprefs.xml file.

Procedure

1. Open the following file for editing: [VSC_home]\etc\kamino\kaminoprefs.xml, where [VSC_home] is your VSC for VMware vSphere installation directory.
2. Modify the following entry by replacing log4j.properties with the new log configuration filename:

```
<entry key="log4j.config.file" value="log4j.properties" />
```
3. Place the new log configuration file in the following directory:
[VSC_home]\etc\kamino
4. Restart the VSC for VMware vSphere service.

Note: The VSC for VMware vSphere service must be restarted only if the kaminoprefs.xml is modified to point to a different configuration file. Changes in the configuration file itself do not require a VSC for VMware vSphere service restart.

Export Files

The Provisioning and Cloning capability creates a network configuration file with information about each virtual machine created. When requested, the Provisioning and Cloning capability creates Citrix XenDesktop export files.

Network configuration file

After each cloning procedure, the Provisioning and Cloning capability creates a .csv network configuration file.

The network configuration file is created in the following location, where [VSC_home] is your VSC for VMware vSphere installation directory:

```
[VSC_home]\etc\kamino\exports
```

The configuration file is written in the following format:

import_generic_date_time.csv

For example: import_generic_2009_03_06_11_04.csv

The file includes the following information for each virtual machine:

- VM Name
- UUID, VMX Path
- Number of CPUs
- Memory in MB
- Guest Full Name
- Guest Alternate Name
- Datastore Name

The file also includes the following information for up to ten NICs:

- NIC x
- NIC x Network Label
- NIC x Address Type
- Manual - Statically assigned MAC address.
- Generated - Automatically generated MAC address.
- Assigned - MAC address assigned by VirtualCenter.
- NIC x WOL Enabled
- Wake-on-LAN enabled or disabled on this virtual network adapter
- NIC x MAC Address

XenDesktop export file

When XenDesktop connection broker output is requested in the cloning wizard, the Provisioning and Cloning capability creates a .csv file in the exports directory.

The XenDesktop connection broker output file is created in the following location, where [VSC_home] is your VSC for VMware vSphere installation directory:

[VSC_home]\etc\kamino\exports

The filename is in the following format:

xenDesktop_timestamp.csv

The file includes the following information:

[ADComputerAccount],[AssignedUser],[VirtualMachine],[HostID]

Provisioning and Cloning vCenter privileges

The following table presents the privileges required to use the Provisioning and Cloning capabilities in vCenter.

Privilege	Actions allowed
Datastore	<ul style="list-style-type: none"> • Allocate space • Browse datastore • Low level file operations • Remove file • Rename datastore
Distributed virtual port group	Modify
Distributed virtual switch	<ul style="list-style-type: none"> • Modify • Port configuration operation • Port setting operation
Extension	<ul style="list-style-type: none"> • Register extension • Update extension
Global	<ul style="list-style-type: none"> • Cancel task • Licenses • Log event • Manage custom attributes • Settings
Host Configuration	<ul style="list-style-type: none"> • Advanced settings • Security profile and firewall • Storage partition configuration
Virtual Storage Console for VMware vSphere	Provisioning and Cloning <ul style="list-style-type: none"> • Configure • Create Rapid Clones • Datastore <ul style="list-style-type: none"> – Manage datastores – Provision • Redeploy clones • Reclaim space • Distribute templates
Network	Assign network
Resource	Assign virtual machine to resource pool
Tasks	<ul style="list-style-type: none"> • Create tasks • Update tasks

Privilege	Actions allowed
Virtual Machine Configuration	<ul style="list-style-type: none"> • Add existing disk • Add new disk • Add or remove device • Advanced • Change CPU count • Change resource • Disk change tracking • Extend virtual disk • Host USB device • Memory • Modify device settings • Raw device • Remove disk • Rename • Settings • Swapfile placement • Upgrade virtual hardware
Virtual Machine Interaction	<ul style="list-style-type: none"> • Answer question • Configure CD media • Configure floppy media • Device connection • Power off • Power on
Virtual Machine Inventory	<ul style="list-style-type: none"> • Create from existing • Create new • Remove • Unregister
Virtual Machine Provisioning	<ul style="list-style-type: none"> • Allow disk access • Allow read-only disk access • Clone template • Clone virtual machine • Create template from virtual machine • Customize • Deploy template • Read customization specifications

Optimizing and migrating datastores and virtual machines

The Optimization and Migration capability provides a simple interface for performing online alignments and migrations of virtual machines. You can align VMFS datastores without having to power down the virtual machine. In addition, you can review the alignment status of the virtual machines and migrate groups of virtual machines into new or existing datastores.

Having the virtual machines aligned can improve performance.

The interface for the Optimization and Migration capability is integrated with VMware vCenter Server and works in conjunction with the VMware Storage vMotion feature to enable a running virtual machine to be moved between datastores.

The Optimization and Migration capability allows you to perform the following tasks:

- Online alignment

The online alignment tool lets you align virtual machines without having to power them down.

- Datastore scans

By scanning datastores, you can determine which virtual machines are misaligned. You can scan all the datastores at the same time, or only the datastores you select.

- Scheduled scans

You can schedule scans of the datastores so that they occur automatically at specific times.

Note: It is a good practice to schedule the scans to occur during non-critical production times. The time required to perform a scan can increase as more virtual machines are scanned.

- Migration of virtual machines

The migration feature allows you to migrate virtual machines either singly or as a group.

Note: Migrating multiple virtual machines at one time is very I/O intensive. Your system is likely to slow down while the migration is in process. You might want to limit the number of virtual machines that you migrate at one time to avoid over-stressing your system during the migration.

- Sorts of virtual machines

The sort feature places the virtual machines in folders that let you quickly determine their states and the actions you can take. Based on the folders, you know the following:

- Whether a virtual machine is functionally or actually aligned.

The **Aligned > Functionally Aligned** and the **Aligned > Actually Aligned** folders contain these virtual machines.

- If a misaligned virtual machine can be aligned using the online feature of the Optimization and Migration capability.

The **Misaligned > Online Migration** folder contains these virtual machines.

- If a misaligned virtual machine must be aligned using an offline tool, such as VMware vCenter Converter.

To use an offline alignment tool, you must power off the virtual machine.

The **Misaligned > Offline Migration** folder contains these virtual machines.

- If a virtual machine is not aligned and cannot be aligned.

The Optimization and Migration capability places the virtual machine in the **Misaligned > Other** folder if it cannot be aligned. This can happen if the virtual machine

- Has disks that are inaccessible
- Has a disk size of 0
- Has more than one disk with different offsets
- Has multiple disks spanning multiple datastores
- Does not have any partitions
- Is an independent disk
- Reports an error during read attempts
- Is a dynamic disk

Note: If you use the Optimization and Migration capability when you have a dynamic disk, the capability might give you a false indication of alignment.

Types of alignments

You can use the Optimization and Migration capability to perform online alignments, and to sort virtual machines so that you can see which ones require you to use another tool and perform an offline alignment. The type of alignment you perform determines whether the virtual machine is functionally aligned or actually aligned.

Online and offline alignments

When you use the Optimization and Migration capability to perform an online alignment, you do not need to power down a virtual machine. If the Optimization and Migration capability is not able to perform an online alignment, you must power down the virtual machine and use a tool such as the VMware vCenter Converter to align the virtual machines offline. The Optimization and Migration capability sorts the virtual machines into folders based on whether you can perform an online alignment, must perform an offline alignment, or cannot align the virtual machine.

Note: At this time, you can only use the Optimization and Migration capability to provide online alignments of VMFS datastores.

During an online alignment, the Optimization and Migration capability moves the virtual machine you specify to a single, optimized VMFS datastore and performs a functional alignment. The datastore can be either a new or existing datastore.

If you have multiple virtual machines that have the same misalignment, you can migrate them in a batch. The Optimization and Migration capability moves all the virtual machines to the same, optimized, VMFS datastore.

Note: Migrating multiple virtual machines at one time is very I/O intensive. Your system is likely to slow down during the migration. You may want to limit the number of virtual machines that you migrate at one time to avoid over-stressing your system during the migration.

Functional and actual alignments

Depending on whether you use the Optimization and Migration capability to perform an online alignment or another tool to perform an offline alignment, your virtual machines are aligned either functionally or actually.

In a functional alignment, the Optimization and Migration capability moves the misaligned virtual machines to an optimized datastore that uses a prefix to ensure that the virtual machines align on correct boundaries. As a result, the virtual machines perform as though they are aligned.

Note: If you clone a virtual machine that has been functionally aligned, the clone will be misaligned. When you use the Provisioning and Cloning capability to clone virtual machines, it warns you if the source is misaligned. Other cloning tools do not provide a warning. As a result, it is recommended that you use only the Provisioning and Cloning capability to clone datastores.

In an actual alignment, the partitions of the virtual machine's hard disk are aligned to the storage system and start at the correct offset. For this type of alignment, you perform an offline alignment, which modifies the contents of a virtual hard disk to align the virtual machine.

The Optimization and Migration capability does not modify the contents of a virtual hard disk.

Important notes about using the Optimization and Migration capability

When you use the Optimization and Migration capability, it is important to understand some key points about how it works.

The following points can affect the results you see when you use the Optimization and Migration capability:

- Multiple partitions on a disk and different offsets for the partitions

When a virtual machine has multiple partitions on a disk and the partitions have different offsets, the Optimization and Migration capability functionally aligns the virtual machines to the offset of the largest partition

In such cases, the Optimization and Migration capability migrates the virtual machine to the datastore offset of the largest partition.

If you have a smaller partition that is accessed more frequently than the larger partition, you might not see an improvement in performance when the larger, less active partition is aligned.

- Multiple disks with different offsets

If a virtual machine has more than one disk with different offsets, you cannot use the Optimization and Migration capability to align that virtual machine.

You can use the capability to migrate that virtual machine to another datastore. You can also power down the virtual machine and use an offline alignment tool such as VMware vCenter Converter to perform an alignment

The Optimization and Migration capability places these virtual machines in the **Misaligned > Offline Migration** folder in the Virtual Machine Alignment panel.

The Optimization and Migration capability can only align virtual machines listed in the **Misaligned > Online Migration** folder.

- Windows 2008 R2 SP1

The Optimization and Migration capability cannot scan a virtual machine running Windows 2008 R2 SP1.

This is a restriction of VMware's Virtual Disk Development Kit (VDDK). In most cases, this is not a problem because these disks are normally aligned.

In the Virtual Machine Alignment panel, the Optimization and Migration capability lists these virtual machines in the **Misaligned > Other** folder.

- VMware's Storage Distributed Resource Scheduler (SDRS) feature

If you want to use VMware's SDRS feature with datastores created by the Optimization and Migration capability, you should not put datastores with varying offsets in a single datastore cluster.

In addition, do not mix optimized and non-optimized datastores in the same datastore cluster.

- Lock management feature

The lock management feature of VSC for VMware vSphere ensures that when one capability is using a target datastore or virtual machine, other capabilities cannot use that datastore or virtual machine at the same time.

For example, if you are using the Provisioning and Cloning capability to perform a cloning operation, you can use the Optimization and Migration capability to scan and create a datastore. However, you cannot use the Optimization and Migration capability to align the same virtual machine that the Provisioning and Cloning capability is using.

- Migration of multiple virtual machines

You can migrate multiple virtual machines at one time; however, each migration is I/O intensive and can slow down your system while it is underway.

Before you migrate multiple virtual machines, consider what your environment can handle. Because of the I/O load, you might want to limit the number of virtual machines you include in a single migration to avoid over-stressing your system.

- Deduplication temporarily turned off

When you use the Optimization and Migration capability to migrate a virtual machine, deduplication is temporarily turned off. As a result, the migration might require more space.

- VAAI extended copy feature

If a datastore contains virtual machines that have been aligned using the Optimization and Migration capability, you cannot use the VAAI extended copy feature.

The extended copy feature is not available on these optimized datastores. To determine if a datastore has been optimized, check the Scan Manager panel to see if there is a **yes** in the Optimized column.

- VMware Storage vMotion

If you use VMware Storage vMotion, you must make sure that you specify the correct optimized datastore as the destination.

Performance problems can occur if you specify the wrong datastore.

Whenever you migrate misaligned virtual machines to optimized datastores using VMware Storage vMotion, you should perform the task from within the Optimization and Migration capability.

The Optimization and Migration workflow

The Optimization and Migration capability lets you scan datastores and correct the alignment of certain misaligned virtual machines (VM) without having to power down the VM.

The following high-level workflow takes advantage of the Optimization and Migration capability features.

Note: The sections that follow provide detailed information on these steps.

1. Go to the Scan Manager panel and click **Scan now** to scan the datastores and VMs.

This action provides information about the datastores and whether they have been optimized.

You can specify which datastores you want to include in the scan. Depending on how many datastores are scanned, it may take a few minutes before the Virtual Machine Alignment panel folders are populated.

2. Go to the Virtual Machine Alignment panel to view the alignment status of the VMs.
3. Check which misaligned VMs can be corrected with the online alignment tool.
4. Use the Optimization and Migration capability to migrate those VMs to a new or existing datastore to correct the alignment issue.
5. Alternatively, select VMs that you wish to migrate.

You can migrate the VMs as a group as long as the VMs have the same misalignment. When you migrate VMs, the Virtual machine migration wizard starts. This wizard provides a series of dialog boxes that prompt you for the information you need to provide.

6. Maintain an optimized environment by scanning the datastores on a regular basis.

Scanning the datastores

You can use the Optimization and Migration capability to scan VMFS datastores and see which ones are optimized.

About this task

As part of the scan process, the Optimization and Migration capability takes VMware snapshots of the virtual machines that are powered on. As soon as the snapshots are no longer needed, it quickly deletes them to free up that space.

Procedure

1. Select the Optimization and Migration capability Scan Manager panel.
2. Specify which datastores you want the Optimization and Migration capability to scan. You have the following options:
 - Click **Scan now** without selecting anything.
 - Exclude certain datastores from being scanned.
Check the box next to the names of those datastores and click **Exclude**.
 - Select specific datastores for scanning.
Check the box next to the names of those datastores.
 - Include datastores that previously were excluded.

- Check the box next to their names and click **Include**.
3. Click either **Scan selected** or **Scan now**. When you choose **Scan selected**, the Optimization and Migration capability scans the datastores with checks next to their names. If you have not selected any datastores, choose the **Scan now** option, which scans the entire datacenter.

After the scan completes, the Optimization and Migration capability displays information about the datastores, including the following:

- Whether the datastore is optimized
- Whether it was included in the most recent scan
- When the last scan of the datastore occurred

What to do next

After you have scanned the datastores, you can go to the Virtual Machine Alignment panel and check the alignment status of the virtual machines.

Scheduling a scan of datastores

The Optimization and Migration capability lets you set up automatic scans of the datastores.

Procedure

1. Select the Optimization and Migration capability Scan Manager panel.
2. Specify which datastores you want the Optimization and Migration capability to scan. You have the following options:
 - Exclude certain datastores from being scanned.
Check the box next to the names of those datastores and click **Exclude**.
 - Select specific datastores for scanning.
Check the box next to the names of those datastores.
 - Include datastores that previously were excluded.
Check the box next to their names and click **Include**.
3. Click **View/edit schedule**.
4. In the pop-up box that appears, specify the day and time when you want the scan to occur. You can specify multiple days. All scans will occur at the same time on each day that you specify.
If you do not use the **Exclude** button, the Optimization and Migration capability scans the all of the datastores.

Note: It is a good practice to schedule the scans to occur during non-critical production times. When scanning virtual machines, the Optimization and Migration capability uses VMware snapshots. Taking these snapshots increases the time required for the scan.
5. When you have the schedule for scans set up, click **Save**. The Optimization and Migration capability automatically scans the datastores according to the schedule you set up.

Performing an online alignment

In many cases, the Optimization and Migration capability allows you to perform online alignments of virtual machines. As a result, you do not have to power the virtual machines down during the alignment procedure. The Optimization and Migration capability aligns the virtual machines by using the Virtual machine migration wizard to migrate them into a single, optimized VMFS datastore.

Before you begin

You should perform a full scan of the datastores in order to populate the database with information about which virtual machines need to be aligned.

About this task

The following are some points to keep in mind about this task:

- In some cases, virtual machines cannot be aligned using the Optimization and Migration capability's online tool.
When a virtual machine cannot use online alignment feature, the Optimization and Migration capability places the virtual machine in the **Offline Migration** folder. You must power down the virtual machine and use an offline tool such as the VMware vCenter Converter to align it.
- If you are working with a Storage Distributed Resource Scheduler (DRS), you must create a storage DRS cluster and make sure it is available for specific offset.
- The VAAI extended copy feature is not available on optimized datastores.
- You should also make sure that you do not mix optimized datastores with non-optimized datastores.
- You cannot cancel an online alignment that is in progress.
- The Optimization and Migration capability only supports online alignment of VMFS datastores.
- In some cases, if the virtual machine is being used by another capability, you cannot perform an online alignment.

Procedure

1. Select the Optimization and Migration capability Virtual Machine Alignment panel.
2. Open **Misaligned > Online Migration**.
3. Select the virtual machines that you want to align. You have the following options:
 - To align all the listed virtual machines, select **Migrate all**.
 - To align specific virtual machines, check the box next to the names of the virtual machines you want to align and then click **Migrate**.
4. The Virtual machine migration wizard launches and prompts you for information about the migration.
 - a. Select the target storage controller that will contain the virtual machines.
 - b. If you are using vFiler units, select the **Set vFiler Context** check box and select a unit from the drop-down list.
5. Specify whether you want to migrate the virtual machines to an existing datastore or create a new datastore.

6. If you are using an existing datastore, select the datastore from the list of datastores displayed in the Choose a datastore for the virtual machines dialog box.

This dialog displays information about the existing datastores, including the following:

- **Datastore name**
- **Datastore type**

Note: The Optimization and Migration capability only supports VMFS datastores.

- **Capacity** of the datastore in GBs
 - **Free space** available on the datastore in GBs
7. If you are creating a datastore, specify its type: **VMFS** or **NFS**.

Note: This release of the Optimization and Migration capability only supports online alignment of VMFS datastores.

To continue without creating a new datastore, click **Next**.

8. In the Specify the details for the new datastore dialog box, provide details about the datastore you want to create.

- **Protocol** (VMFS only) - FCP or iSCSI.
- **Size (GB)** - Specify the maximum size for the datastore.

The maximum size depends on the controller and space available. For details, see the *Data ONTAP Storage Management Guide* for your Data ONTAP release.

- **Datastore name** - Supply a name for the datastore you are creating.
 - **Create new volume container** (VMFS only) - Create a volume with the same name as the LUN. If a volume with that name already exists, the Optimization and Migration capability appends a number to the volume name; for example, **Volname01**.
 - **Volume** (VMFS only) - Select a volume from drop-down list of available volumes.
 - **Thin provision** - Select this option to set space reserve to none and disable space checks.
 - **Block size** (VMFS only) - Select the block sizes. For VMFS-5, block size is fixed at 1 MB.
9. Review the summary page and click **Apply** to proceed. To return to previous pages and modify settings, click **Back**. The Optimization and Migration capability displays the status of the virtual machines as they are aligned.
 10. After the online alignment completes, the virtual machines you just aligned no longer appear in the **Misaligned > Online Migration** folder.

Migrating virtual machines

You can use the Optimization and Migration capability to migrate a group of virtual machines (VM) at one time.

About this task

Migrating multiple VMs at one time is very I/O intensive. Your system is likely to slow down while the migration is in progress. You may want to limit the number of VMs that you migrate at one time to avoid over-stressing your system during the migration.

You cannot cancel a migration that is in progress.

Migration is not restricted to optimized datastores. You can select groups of VMs to migrate to any datastore as long as the source and destination datastore have the same offset.

Note: In some cases, when the VM is being used by another capability, you cannot use the migration features.

The Optimization and Migration capability only supports online alignment of VMFS datastores.

Procedure

1. From the Virtual Machine Alignment panel, select the VMs that you want to migrate. Click **Migrate all** to migrate all the VMs or **Migrate** to migrate only the VMs you have checked.
2. The Virtual machine migration wizard launches and prompts you for information about the migration.
 - a. Select the target storage controller that will contain the VMs.
 - b. If you are using vFiler units, select the **Set vFiler Context** check box and select a unit from the drop-down list.
3. Specify whether you want to migrate the VMs to an existing datastore or create a new datastore.
4. If you are using an existing datastore, select the datastore from the list of datastores displayed in the Choose a datastore for the virtual machines dialog box.

This dialog displays information about the existing datastores, including the following:

- **Datastore name**
 - **Datastore type**
 - **Capacity** of the datastore in GBs
 - **Free space** available on the datastore in GBs
5. If you are creating a datastore, specify its type: **VMFS** or **NFS**. To continue without creating a new datastore, click **Next**.

Note: This release of the Optimization and Migration capability only supports online alignment of VMFS datastores.

6. When you create a new datastore, the Specify the details for the new datastore dialog box prompts you for details.
 - **Protocol** (VMFS only) - FCP or iSCSI.
 - **Size (GB)** - Specify the maximum size for the datastore.
The maximum size depends on the controller and space available. For details, see the *Data ONTAP Storage Management Guide* for your Data ONTAP release.
 - **Datastore name** - Supply a name for the datastore you are creating.
 - **Create new volume container** (VMFS only) - Create a volume with the same name as the LUN. If a volume with that name already exists, the Optimization and Migration capability appends a number to the volume name; for example, **Volname01**.
 - **Volume** (VMFS only) - Select a volume from drop-down list of available volumes.

- **Thin provision** - Select this option to set space reserve to none and disable space checks.
Selecting thin provisioning for a datastore can let you over subscribe the storage controller. The limit for resizing the datastore is 10 times its initial size.
 - **Block size** (VMFS only) - Select block sizes. For VMFS-5, block size is fixed at 1 MB.
7. Review the summary page and click **Apply** to proceed. To return to previous pages and modify settings, click **Back**.

Backing up and restoring data

The Backup and Recovery capability of Virtual Storage Console for VMware vSphere enables you to rapidly back up and recover multihost configurations running on IBM N series storage systems.

You can use this capability to perform the following tasks:

- Create on-demand backups of individual virtual machines, datastores, or a datacenter
- Schedule automated backups of individual virtual machines, datastores, or a datacenter
- Support virtual machines and datastores that are located on either NFS directories or VMFS file systems
- Mount a backup to verify its content prior to restoring it
- Restore datastores to their original locations
- Restore virtual machines to their current locations
- Restore virtual machine disks (VMDKs) to their current locations or alternate locations

Related concepts:

“Configuring the Backup and Recovery capability” on page 98

“Managing backups” on page 100

“Restoring data from backups” on page 104

“Single file restore” on page 108

“VSC CLI commands” on page 116

Related reference:

“Resolution of issues with the Backup and Recovery capability” on page 170

Backup and Recovery requirements

Your datastore and virtual machines must meet Snapshot and SnapRestore requirements before you can use the Backup and Recovery capability.

Following are some Snapshot and SnapRestore requirements:

- Snapshot protection is enabled in the volumes where those datastore and virtual machine system images reside.
- SnapRestore technology is licensed for the storage systems where those datastore and virtual machine system images reside.

Backup and Recovery requirements for optional SnapMirror protection

Optional SnapMirror protection requires that you perform certain configuration tasks and establish relationships between virtual objects in your environment prior to using the Backup and Recovery capability.

Following are some SnapMirror requirements:

- The volumes containing the active datastore and virtual machine images must be configured as SnapMirror source volumes.

- The source volumes must have a SnapMirror relationship with target volumes on a second SnapMirror-licensed storage system that is located a safe distance from the source storage system.
- The host names and IP addresses of the SnapMirror source and destination storage systems must be resolvable for the SnapManager for Virtual Infrastructure server, either through a configured DNS server or through host entries added to the host file on the SnapManager for Virtual Infrastructure server.
- Cluster or Vserver administrators must create a Vserver management LIF, which is required by the Backup and Recovery capability to update SnapMirror relationships for storage systems operating in Cluster-Mode.

Configuring the Backup and Recovery capability

Before you begin using the Backup and Recovery capability to schedule backups and restore your datastores, virtual machines, or virtual disk files, you must add the storage systems that contain the datastores and virtual machines for which you are creating backups. You can also create a custom user account for a storage system.

Authentication methods in the Backup and Recovery capability

You can override the default authentication method by creating a custom user account through the CLI, which enables you to log in with authentication credentials other than your Windows credentials.

Custom user accounts for accessing a storage system

A non-root or non-administrator account might be required for the Backup and Recovery capability to access a specific storage system.

In this case, you must create a custom storage system account with a new storage system role, group, and user as described in the following table.

Item	Description
Role	The new role must allow the Backup and Recovery capability to access the storage system data through its APIs.
Group	A storage system maintains groups as a collection of roles. The group you create must contain your new role.
User	A user account that the Backup and Recovery capability uses to access a storage system must be a member of a group that contains a role. You can create this user and assign a password to it, then you are able to add a storage system to Backup and Recovery with the assigned user name and password.

For more information on how to manage users on your storage system, refer to your storage system's administrator guide.

Creating custom users

You can use the CLI to create a custom user. The credentials of a custom user provide the same access to commands and features as an administrator who logs in using the default Windows credentials authentication method.

Procedure

1. Double-click the **VSC CLI** desktop icon or navigate to **Start > All Programs > IBM > Virtual Storage Console > IBM N series VSC CLI**.
2. Enter the **smvi servercredential set** command.
3. Enter a user name and a password.

Creating a custom storage system user account

You can use your storage system's CLI to create a custom storage system account with a new storage system role, group, and user.

About this task

Complete the following steps from the CLI of the storage system that the Backup and Recovery capability needs to access.

Procedure

1. Use the following command to create a role named `api-access` with the minimum configuration required for the Backup and Recovery capability to access the storage system.

```
useradmin role add api-access -a api-*,login-http-admin,cli-ifconfig
```
2. Use the following command to create a group named `api-group` which contains the `api-access` role.

```
useradmin group add api-group -r api-access
```
3. Use the following command to create a user named `smvi-user` as a member of the `api-group` group.

```
useradmin user add smvi-user -g api-group
```
4. To set the user's password, run the **passwd** command as root. The storage system prompts you for the account name that you want to change, followed by the new password for this account.

How the Backup and Recovery capability discovers vFiler units

The Backup and Recovery capability supports vFiler tunneling on the physical storage system that contains the vFiler units. In certain environments, the Backup and Recovery capability can access only the physical storage system, not the vFiler unit, for communication on a storage network.

With vFiler tunneling enabled, you can use the Backup and Recovery capability to create Snapshot copies. When you add a physical storage system to the Backup and Recovery capability, it automatically discovers all the associated vFiler units. You do not need to manually add the vFiler units.

Related tasks:

“Enabling discovery and management of vFiler units” on page 22

“Enabling discovery and management of vFiler units on private networks” on page 23

Managing backups

You can perform on-demand backups or schedule automated backups of individual virtual machines, datastores, or a datacenter using the Backup and Recovery capability.

After you add a new backup job using the Backup wizard, you can set a schedule and specify a retention policy for the backup. You can also change the schedule and retention policy for a backup and delete or suspend a backup job.

Considerations for adding a backup job

Before you use the Backup wizard to add, schedule, and automate backups, you should be aware of the information that you can specify to ensure that the backup schedule, the backup retention policy, and the email alerts of backup activity for your job perform as expected.

When you add a new backup job, you can specify whether you want to initiate a SnapMirror update on the virtual entities that are to be backed up or create VMware snapshots for every backup. If you select virtual entities that span multiple datastores, you can specify one or more of the spanning datastores to be excluded from the backup.

If you want to run a backup script that is installed on the server with this job, you can choose the scripts that you want to use. You can specify the hourly, daily, weekly, or monthly schedule that you want applied to your backup job or you can add a backup job without attaching a schedule to the backup. You can also specify the maximum number of days and maximum number of backups and the e-mail alerts for this backup job.

Note: If you create a pre-backup or post-backup script that results in an output file when the script is run, the output file is saved to the same directory where you initially installed the pre-backup or post-backup script.

Backing up a virtual machine

You can use the Backup wizard to add a new backup job for a virtual machine. You can schedule and automate your backups, specify a retention policy, and create an automated policy for email alerts.

Before you begin

The vSphere Client must be connected to a vCenter Server to create backups.

About this task

The Backup and Recovery capability uses the Monitoring and Host Configuration storage system discovery information. You can add a new storage system to the Backup and Recovery capability by specifying the host name or IP address of the storage system in the Monitoring and Host Configuration capability.

During backup or recovery of a virtual machine, the Backup and Recovery capability does not allow other backup or recovery operations on that virtual machine to start. The Backup and Recovery capability delays any other backup or recovery operations until the current backup or recovery is finished.

Procedure

1. Select a virtual machine in the Inventory panel, and then right-click the node and select **IBM N series > Backup and Recovery > Schedule a Backup**. The Backup wizard appears.
2. Type a backup job name and description and click **Next**.
3. Select one of the following options:
 - If you want to start a SnapMirror update on the selected entities concurrent with every backup, select **Initiate SnapMirror update**.
For this option to execute successfully, the selected entities must reside in volumes that are already completely configured as SnapMirror source volumes.
The SnapManager for Virtual Infrastructure server should be able to resolve the host name and IP address of the source and destination storage systems in the `snapmirror.conf` file.
 - If you want to create a VMware snapshot for each backup, select **Perform VMware consistency snapshot**.
 - If you want to include independent disks from datastores that contain temporary data, select **Include independent disks**.
4. Select the virtual entities available for this backup job and click **Next**.
5. Select one or more backup scripts and click **Next**. If an error message appears indicating that at least one of the selected scripts has been deleted, you can save the backup job without any script in the selected scripts list, thereby removing the deleted script from the job. Otherwise, the backup job continues to use the deleted script.
6. Select the hourly, daily, weekly, or monthly schedule that you want for this backup job and click **Next**.
7. Select **One time only** and click **Delete this job after backup is created** if you do not want to retain this backup job.
8. Type the user name and password for the vCenter Server and click **Next**.
9. Select the maximum number of days or the maximum number of backups and click **Next**.
10. Select the email alerts and click **Next**. You can add multiple email addresses by using semicolons to separate each email address.
11. Review the summary page and click **Finish**.
Select the **Run Job Now** check box to immediately run the job.

Backing up a datastore or datacenter

You can use the Backup wizard to add a new backup job for an entire datacenter, a datastore, a particular set of datastores, or a particular set of virtual machines. You can schedule and automate your backups, specify a retention policy, and create an automated policy for email alerts.

Before you begin

The vSphere Client must be connected to vCenter Server to create backups.

About this task

The Backup and Recovery capability uses the Monitoring and Host Configuration storage system discovery information. You can add a new storage system to the Backup and Recovery capability by specifying the host name or IP address of the storage system in the Monitoring and Host Configuration capability.

During backup or recovery of a virtual machine, the Backup and Recovery capability does not allow other backups or recoveries of that virtual machine to start. The Backup and Recovery capability delays any other backup or recovery operations until a current backup or recovery operation is finished.

Procedure

1. Select a datacenter or datastore in the Inventory panel, and then right-click the node and select **IBM N series > Backup and Recovery > Schedule a Backup**. The Backup wizard appears.
2. Type a backup job name and description and click **Next**.
3. Select one of the following options:
 - If you want to start a SnapMirror update on the selected entities concurrent with every backup, select **Initiate SnapMirror update**.
For this option to execute successfully, the selected entities must reside in volumes that are already completely configured as SnapMirror source volumes.
The SnapManager for Virtual Infrastructure server should be able to resolve the host name and IP address of the source and destination storage systems in the `snapmirror.conf` file.
 - If you want to create a VMware snapshot for each backup, select **Perform VMware consistency snapshot**.
 - If you want to include independent disks from datastores that contain temporary data, select **Include independent disks**.
4. Select the virtual entities available for this backup job and click **Next**.
5. Select one or more backup scripts and click **Next**. If an error message appears, indicating that at least one of the selected scripts has been deleted, you can save the backup job without any script in the selected scripts list, thereby removing the deleted script from the job. Otherwise, the backup job continues to use the deleted script.
6. Select the hourly, daily, weekly, or monthly schedule that you want for this backup job and click **Next**.
7. Select **One time only** and click **Delete this job after backup is created** if you do not want to retain this backup job.
8. Type the user name and password for the vCenter Server and click **Next**.
9. Select the maximum number of days or the maximum number of backups to retain and click **Next**.
10. Select the email alerts and click **Next**. You can add multiple email addresses by using semicolons to separate each email address.
11. Review the summary page and click **Finish**.
Select the **Run Job Now** check box to immediately run the job.

Starting a one-time backup

You can perform a one-time backup on a virtual machine or datastore. This type of backup is useful if you do not want to schedule regular backups for a particular virtual machine or datastore.

Procedure

1. Select a virtual machine or datastore in the Inventory panel, and then right-click the node and select **IBM N series > Backup and Recovery > Backup Now**. The Backup Now window appears.
2. Clear the **Automatically name backup** check box if you do not want the backup name to be automatically generated.
3. Type a backup name in the Name text box and click **Backup Now**.
4. Select one of the following options:

- If you want to start a SnapMirror update on the selected entities concurrent with every backup, select **Initiate SnapMirror update**.

For this option to execute successfully, the selected entities must reside in volumes that are already completely configured as SnapMirror source volumes.

The SnapManager for Virtual Infrastructure server should be able to resolve the host name and IP address of the source and destination storage systems in the `snapmirror.conf` file.

- If you want to create a VMware snapshot for each backup, select **Perform VMware consistency snapshot**.
- If you want to include independent disks from datastores that contain temporary data, select **Include independent disks**.

Editing a backup job

You can use the Job Properties dialog box to modify the name and description, the datastores and virtual machines that are assigned, the backup scripts, the user credentials, the schedule, the retention policy, or the e-mail alerts for an existing backup job.

Procedure

1. Select a virtual machine or datastore in the Inventory panel, and then click the IBM N series tab.
2. Click **Backup** under Backup and Recovery in the navigation pane.
3. Select the backup job whose properties you want to modify.
4. Click **Edit**, and then click the appropriate tab for the properties that you want to modify for this backup job.
5. Modify backup job properties as necessary and click **OK** to change the properties.

Deleting a scheduled backup job

You can select and delete one or more backup jobs from the list of scheduled jobs, but you cannot delete any backup jobs that are running.

Procedure

1. Select a virtual machine or datastore in the Inventory panel and click the IBM N series tab.
2. Click **Backup** under Backup and Recovery in the navigation pane.

3. Select one or more backup jobs that you want to delete. In the Entities page, note the existing datastore and virtual machines currently associated with the selected backup job. When you delete the selected backup job, its backup operations are no longer performed on these entities.
4. Click **Delete**, and then click **Yes** at the confirmation prompt.

Suspending an active backup job

You can suspend an active backup job and its scheduled operations without deleting the job.

Procedure

1. Select a virtual machine or datastore in the Inventory panel, and then click the IBM N series tab.
2. Click **Backup** under Backup and Recovery in the navigation pane.
3. Select the active backup job that you want to suspend. In the Entities page, note the existing datastore and virtual machines currently associated with the selected backup job. When you suspend the selected backup job, its backup operations are no longer carried out on these entities.
4. Right-click the selected backup job, and then select **Suspend**.
5. Click **Yes** on the confirmation prompt.

Resuming a suspended backup job

You can resume and run a suspended backup job.

Procedure

1. Select a virtual machine or datastore in the Inventory panel, and then click the IBM N series tab.
2. Click **Backup** under Backup and Recovery in the navigation pane.
3. Select the suspended backup job that you want to resume.
4. Right-click the selected backup job, and then select **Resume**.

Note: The **Resume** option is not active unless the selected backup job is in a suspended state.

5. Click **Yes** on the confirmation prompt.

Restoring data from backups

You can use the Backup and Recovery capability to restore your virtual machines and datastores from backups that contain them. VSC for VMware vSphere supports restoring from local backups and from backups that contain VMware snapshots.

The Backup and Recovery capability uses FlexClone technology to restore data on systems running Data ONTAP 7.x and, in conjunction with **snap restore hidden flag**, Data ONTAP 8.1.

For VMware Virtual Machine File System (VMFS) datastores, if FlexClone is not licensed, the Backup and Recovery capability uses LUN clone technology on systems running Data ONTAP 7.x and 8.1.x. For Network File System (NFS) datastores in cluster configurations, the Backup and Recovery capability uses **snap restore hidden flag**, while VMFS datastores use SIS clone technology.

Where to restore a backup

VSC for VMware vSphere enables you to select the original (default), the current, or an alternate location as the restore destination by using the Restore wizard.

- Original location (default)

The backup copy of an entire datastore is restored to the original location. You set this location by choosing an entire datastore. This type of restore operation is referred to as an “in-place restore” operation.

- Current location

The backup copy of a single virtual machine, or a copy of a virtual machine's disk files, is restored to a new, current location after it has been migrated from its original location. You set this location by choosing an entire virtual machine.

- Alternate location

The backup copy of the virtual machine disk files is restored to an alternate location. This type of restore operation is referred to as an “out-of-place restore” operation. You set this location by selecting the destination location as a different datastore from the one on which the virtual disk currently resides by setting the **Particular virtual disks** option.

Restore operations using data that was backed up with failed VMware consistency snapshots

Even if a VMware consistency snapshot for a virtual machine fails, the virtual machine is nevertheless backed up. You can view the backup copy in the Restore panel and use it for restore operations.

During backup operations, a VMware quiesced snapshot is created by default. When creating a snapshot, the virtual machine pauses all running processes on the guest operating system so that file system contents are in a known consistent state when the snapshot is taken. Despite the VMware snapshot failure, the virtual machine is still included in the Snapshot copy. You can restore a virtual machine without manually removing it from a Snapshot copy.

You can restore a potentially functional virtual machine despite failing the VMware process of creating a quiesced snapshot, and view its status during quiesced backup operations in the backed-up entities section of the Restore panel. The Quiesced column can display the following values:

- Yes, if a VMware snapshot operation is successful and the guest operating system was quiesced.
- No, if the VMware snapshot operation failed because the guest operating system could not be quiesced.
- Not Applicable, for entities that are not virtual machines.

Restoring data from backups

You can restore a datastore, virtual machine, or its disk files to its original location or an alternate location. From the Restore panel, you can sort the backup listings by name, date, or other search criteria to help you find your backups.

From the Restore panel, you can do the following:

- Restore a datastore, virtual machine, or its VMDKs from a local backup to an original location
- Restore VMDKs from a local backup to a different location
- Restore from a backup that has a VMware snapshot

Searching for backups

You can locate a specific backup by searching the backup table. After you locate a backup, you can then restore from it, view its status, or delete it.

Procedure

1. From the navigation panel, click Restore panel to display the available backups.
2. In the Search text box, type one or more search terms and press **Enter**. Available criteria are the backup ID, VMware snapshot, mount state, or resource name. A filtered list displays in the Restore panel.

Restoring a datastore

You can use the VSC for VMware vSphere to restore a datastore. By doing so, you overwrite the existing content with the backup you select.

Procedure

1. In the Restore panel, select a backup of the datastore.
2. Click **Restore**. Restoring a datastore powers off all virtual machines.
3. In the **Restore dialog box**, click **Restore**.

Restoring a virtual machine or its VMDKs

You can restore an entire virtual machine or particular virtual disks of a given virtual machine. By doing so, you overwrite the existing content with the backup you select.

Before you begin

You must have already backed up a virtual machine or its VMDKs using before you can restore it.

Procedure

1. Select a virtual machine in the Inventory panel, and then right-click the node and select the **IBM N series > Backup and Recovery > Restore** from the context menu. The Restore wizard opens and lists all backups that include the virtual machine.
2. From the Backups table, select a backup of the virtual machine and click **Next** to display the Virtual Machine Component Selection page.
3. Select one of the following recovery options:

Option	Description
The entire virtual machine	Restores the contents of your virtual machine from a Snapshot copy. The Start VM after restore check box is enabled if you select this option and the virtual machine is registered.
Particular virtual disks	Restores the contents of the VMDKs on a virtual machine. This option is enabled if you uncheck the entire virtual machine option.

4. In the **ESX Host Name field**, select the name of the ESX host. This option is available if you want to restore VMDKs or if the virtual machine is on a VMFS datastore.
5. Click **Next** to display the Summary page, which displays information about the restore operation that is about to be performed.

6. From this page, view the settings and click **Restore** to begin the restore process.

What to do next

After you finish using the Restore wizard and start the restore operation, you can track the progress of the restore operation from the Task tab of the Status panel and monitor the job for possible errors.

Mounting a backup

The Backup and Recovery capability enables you to mount an existing backup onto an ESX server to verify the contents of the backup prior to completing a restore operation or to restore a virtual machine to an alternate location.

About this task

All of the datastores and the virtual machines contained in the backup are mounted to the ESX server that you specify. After it is mounted, the backup copy appears in the vCenter Server user interface.

Procedure

1. Select a virtual machine or datastore in the Inventory panel and click the IBM N series tab.
2. Click **Restore** under Backup and Recovery in the navigation pane.
3. In the Restore panel, select the name of an unmounted backup copy that you want to mount.
4. Click **Mount**.
5. In the Mount Backup dialog box, select the name of the ESX server to which you want to mount the backup. You can mount only one backup each time, and you cannot mount a backup that is already mounted.
6. Click **Mount**.

Unmounting a backup

After you are done verifying the contents of a mounted backup copy, you can unmount it from the ESX server. When you unmount a backup, all of the datastores in that backup copy are unmounted and are no longer visible from the ESX server.

About this task

In some instances, unmounting a backup copy fails if there are virtual entities in use from a previously mounted copy. You must manually clean up the backup prior to remounting it, because its state reverts to not mounted.

In some instances, unmounting a backup copy fails if all of the datastores contained in a backup are in use, but the state of this backup changes to mounted. You can unmount the backup after determining that the datastores are not in use.

Procedure

1. Select a virtual machine or datastore in the Inventory panel and click the IBM N series tab.
2. Click **Restore** under Backup and Recovery in the navigation pane.
3. In the Restore panel, select the name of a mounted backup that you want to unmount.
4. Click **Unmount**.

Note: The backup is unmounted unless the ESX server becomes inactive or restarts during an unmount operation and the job is terminated. In this case, the mount state remains mounted and the backup stays mounted on the ESX server.

5. At the confirmation prompt, click **Yes**.

Single file restore

The Backup and Recovery capability and Restore Agent of Virtual Storage Console provide tools that help you restore backups of your virtual machine disks at the file level.

How these applications coordinate the restore process depends on the network connection between the systems that run Virtual Storage Console and Restore Agent and the level of access that the user has to the destination virtual machine.

How Virtual Storage Console detects network connectivity

Virtual Storage Console assumes that the port group on the VMware vSphere virtual network has a direct connection between the server and the destination virtual machine.

When to manage port group settings

Virtual Storage Console collects information about the port group after it authenticates the user credentials of the vCenter Server. Therefore, to communicate with a Restore Agent instance that is installed on the destination virtual machine, you must check the port group settings on the Single File Restore panel whenever there is a change to the port group settings on the virtual network.

What happens if you do not update the port group

Virtual Storage Console assumes that all port groups can communicate with the SnapManager for Virtual Infrastructure server. If the administrator does not change the virtual network configuration after registering the Restore Agent, then you do not need to manage these settings.

If there are changes to the network configuration after you create a restore session and after you install Restore Agent on the guest virtual machine, in addition to checking the port groups settings, you need to create a new restore session to send the new configuration file to the user to use it with Restore Agent.

The difference between limited and direct connectivity

Users have a varying degree of access to a backup depending on whether there is a connection between the physical and virtual network.

In virtual environments in which the virtual machines are in a separate network from the SnapManager for Virtual Infrastructure server, Restore Agent cannot communicate with Virtual Storage Console. As a result, the user has an *offline* connection and cannot access the list of backups that were made by Virtual Storage Console. If users with limited connectivity want access to these backups, those users must have an administrator create a restore session and send a configuration file to them in an e-mail message.

When there is a direct or *online* connection between the SnapManager for Virtual Infrastructure server and the destination virtual machine, the user can search for mounted backups without administrative help.

By default, all virtual machines have direct connectivity, unless you change the port group setting.

Types of file restore sessions

Backup and Recovery automates the process of restoring single files based on the relationship between the source virtual machine (which was backed up) and the destination virtual machine. Backup and Recovery supports three types of restore sessions.

Self-service

The Backup and Recovery administrator creates a restore session using Backup and Recovery capability. Users can then install Restore Agent on the destination virtual machine, browse the mounted backups on a guest virtual machine, and recover the individual disk file.

Administrator-assisted

This type of file recovery is basically the same as self-service, except that the Backup and Recovery administrator runs Restore Agent and copies the recovered files to a shared location that the user has access to.

Limited self-service

The Backup and Recovery administrator finds the backup copy within a user-specified range of backups and attaches the backed-up disks to the destination virtual machine. The user can then run Restore Agent on a destination virtual machine, browse the mounted backups, and recover the individual disk file.

Manually creating a .sfr file for the Restore Agent

In some environments, security restrictions prevent the email delivery of the Single File Restore (.sfr) file to users. You can use the Backup and Recovery capability to perform a flexible clone copy of a datastore, mount the datastore to an ESX host, and add the virtual disk to a guest machine as a new virtual disk by creating a GuestRestoreClientSession.sfr file, which is delivered to users via email.

About this task

You can manually create this file by extracting a portion of the Backup and Recovery log file.

Procedure

1. After you create a new Single File Restore session, navigate to the Backup and Recovery installation directory that contains a server.log file such as `C:\program files\IBM\SMVI\server\log`.
2. Copy the file to a location where you can access it, such as a shared network drive, and then open the file using a text editor.
3. In the open file, copy all of the text from the string `<standalone>` that corresponds to the date and time when you created the Single File Restore job session, such as `<?xml version=1.0 encoding=UTF-8 standalone=no?>`.
4. Create a new text file named `<filename.sfr>`, and paste the copied text into it. You can use this file to load the configuration for the Restore Agent onto your guest machine.

General configuration settings for single file restore

Using the General tab of the Setup panel, you should set some general configuration parameters so that you can use the single file restore feature.

Setting session defaults

When you set a session default from the Setup panel in the Backup and Recovery capability, any restore session that you subsequently create uses the default settings that you specified.

Procedure

1. In Setup panel, click **Single File Restore**.
2. Click **Edit** and then change the following session defaults:
 - Type the location of the download directory for Restore Agent.
 - Select the number of hours or days before the restore session expires.
The **Default Session Expiration Time** field displays the time in hours.
3. Click **OK**.

Changing the network connection for a port group

You can change the type of restore session for a port group. By doing so, you change the network connection for all the virtual machines on the same subnet.

Procedure

1. In the navigation panel, click **Set up**.
2. In the Setup panel, click **Single File Restore**.
3. Select **Admin Assisted** for the port group.

Results

Backup and Recovery updates the information for the virtual machines that are reported for the port group and in the list of restore sessions in the Single File Restore panel.

Setting the SnapManager for Virtual Infrastructure server address

The restore session configuration file includes the SnapManager for Virtual Infrastructure server IP address and fully qualified domain name, but you might want to change the address or name with a specific one when you have a multi-homed server that has multiple IP addresses.

Procedure

Add the SnapManager for Virtual Infrastructure server IP address to the `/etc/smvi.override` file. Virtual Storage Console will use this value, which can be an IP address or host name, instead of the one previously configured.

Self-service example workflow

To help you understand how a self-service restore session works, imagine that you are a VMware administrator and you need to restore from a backup copy of a virtual machine a critical data file that has become corrupted.

A Virtual Storage Console administrator responds to your ticket and creates a single file restore session. The SnapManager for Virtual Infrastructure server sends

you an e-mail message that provides information you need to access his computer, a link to download Restore Agent application software, and a configuration file.

The configuration file contains data that allows your workstation to communicate with the SnapManager for Virtual Infrastructure server. Rather than the administrator doing the connection work, you can select which disks to connect to and find a backup.

Creating a self-service restore session

You must create a restore session that you or a VMware administrator can use in order to access a guest operating system and find a backup to use.

Before you begin

You must have available the information necessary to use the Add Single File Restore Session wizard effectively:

- The name or IP address of the virtual machine: the backup source and destination
- The e-mail message recipient
- The mount expiration time: 240 hours (or 10 days)

You must be authorized to perform all steps of this task on the vCenter Server. You can configure the authorization credentials from the Setup panel.

Procedure

1. Click the **Single File Restore** link.
2. Click **Add** to start the Add Single File Restore Session wizard.
3. Complete the wizard, using the following values:
 - Source VM Name: VM-XP-EXAMPLE
 - Destination VM Name: VM-XP-EXAMPLE
 - Recipient Email Address(s): JenniferKWhite@example.com
 - Sender Email Address: Jeronim01szewsk@example.com
 - Mount Expiration: 10 days
 - File Restore Access Type: Self-service

Note: In this example, the source virtual machine and the destination virtual machine are the same.

4. Confirm the details of the restore session, and then click **Finish** to complete the wizard.

Results

Your new restore session is listed in the Single File Restore panel.

What to do next

You receive an e-mail notification that contains a link to download Restore Agent, and you install the software.

Installing Restore Agent

After you create a restore session, you receive an email message that provides a link to the Restore Agent installation file and has a restore session configuration attached as a .sfr file. Before you can restore single files on a guest operating system, you must install Restore Agent.

Before you begin

The system upon which you are going to install Restore Agent must have the following software installed:

- Microsoft Management Console 3.0
- Microsoft .NET Framework 3.5 Service Pack 1

In addition, to enable the single file restore feature for an NFS datastore and perform a mount operation, the storage system must have an installed FlexClone license.

For the most current software requirements, see the N series Interoperability Matrices website (accessed and navigated as described in Websites) at www.ibm.com/systems/storage/network/interophome.html.

Procedure

1. Click the link in the email message to download and start the installation process.
2. Follow the displayed instructions.

Load the configuration file

You must load the configuration file that you received in the e-mail message, along with the link to download Restore Agent, onto the destination virtual machine to access the mounted backup virtual machine disk files.

About this task

You can only load the configuration file after you have already installed the Restore Agent on the destination virtual machine.

Procedure

1. Double-click the Restore Agent shortcut icon on your desktop.
2. In the resulting **Load Configuration** window, look for the configuration (.sfr) file.
3. Click **OK**.
4. Choose the virtual machine source to restore from.

Recovering single files from a virtual machine

After you have been given access to files on the destination virtual machine, you can use the Restore Agent window to view the backed up files created by Virtual Storage Console and find the ones that you need to restore from.

Before you begin

When you are restoring data, Restore Agent displays the first-available disk drive letter. If your guest operating system is Windows XP, 7, or 8, you must go to the Windows Disk Management snap-in application to manually assign the disk drive letter to other partitioned space on the disk. For more information about using the

disk management utility, select **Help** from Disk Management.

Procedure

1. From the Restore Agent window, select a file from the **Disk tab**. Alternatively, click the **Backup** tab to find the Snapshot copy by name.
2. Right-click the name of the backup copy and select **Mount**. The **Backup** tab shows the drive letter for each backup copy only if the source virtual machine and destination virtual machine are the same. The contents of the backup copy are written to the new location.

What to do next

After finishing this task, you should clear the configuration cache.

Clear the configuration

After restoring your data from a backed up virtual machine, you should clear the configuration so you can upload another one later.

Procedure

1. From the Action pane in the Restore Agent window, click **Clear Configuration**.
2. In the resulting window, click **OK**. Backup and restore metadata will be removed from Restore Agent and Virtual Storage Console.

Limited self-service example workflow

To help you understand how a limited self-service restore session works, imagine that you are a VMware administrator and you must restore a failed disk on your virtual machine with the help of a Backup and Recovery administrator.

It is Thursday afternoon and one of the files on the virtual machine is corrupt. You must restore the disk over the weekend. You complete a ticket and request access to another workstation so that you can recover the data from a backup copy.

A Virtual Storage Console administrator responds to your ticket and makes the system available from Saturday morning for 48 hours. The SnapManager for Virtual Infrastructure server sends you an email message that provides information you need to log in as an administrator to the Virtual Storage Console computer, a link to a site from which you can download Restore Agent application software, and a configuration file.

The configuration file contains attached disks so that your operating system can see which disks are attached and assign drive letters that enable your workstation to connect to that of the Virtual Storage Console administrator.

Create a limited self-service restore session

The first thing you need to do is create a restore session that a user can use to access a guest operating system restore a virtual machine disk file.

Before you begin

Before creating a restore session, you need to gather the information necessary to complete the Add Single File Restore Session wizard:

- The name or IP address of the virtual machine: the backup source and destination
- The e-mail message recipient and sender

- The mount expiration time: three hours

Ensure that you are authorized to perform all steps of this task on the vCenter Server. You can configure the authorization credentials from the Setup panel.

Procedure

1. Select the **Single File Restore** panel.
2. From the Single File Restore panel, click **Add** to start the Add Single File Restore Session wizard.
3. Complete the wizard, using the following values:
 - Source VM Name: **VM-WXP-EXAMPLE**
 - Destination VM Name: **VM-WXP-EXAMPLE**
 - To Email Address(s): **JordanEKanode@example.com**
 - From Email Address: **LinYaoHuang@example.com**
 - Mount Expiration: **3 days**
 - File Restore Access Type: **Limited Self-Service**

Note: In this example the source virtual machine and the destination virtual machine are the same.

4. Confirm the details of the restore session, then click **Finish** to complete the wizard.

Results

Your new restore session is listed in the Single File Restore panel.

What to do next

You receive an e-mail notification that contains a link to download Restore Agent, and you install the software.

Installing Restore Agent

After you create a restore session, you receive an email message that provides a link to the Restore Agent installation file and has a restore session configuration attached as a .sfr file. Before you can restore single files on a guest operating system, you must install Restore Agent.

Before you begin

The system upon which you are going to install Restore Agent must have the following software installed:

- Microsoft Management Console 3.0
- Microsoft .NET Framework 3.5 Service Pack 1

In addition, to enable the single file restore feature for an NFS datastore and perform a mount operation, the storage system must have an installed FlexClone license.

For the most current software requirements, see the N series Interoperability Matrices website (accessed and navigated as described in Websites) at www.ibm.com/systems/storage/network/interophome.html .

Procedure

1. Click the link in the email message to download and start the installation process.
2. Follow the displayed instructions.

Configuring vCenter Server

After installing Restore Agent, you must configure basic settings for the vCenter Server so that Backup and Recovery can connect to the virtual machine and provide access to its contents.

Before you begin

To use the single file restore feature, vCenter Server must be running ESX 3.5 or later, or vSphere 4.

About this task

You will change the IP address and user credentials for vCenter Server.

Procedure

1. In the navigation pane, click **Set up**.
2. Go to the vCenter Server area of the General tab.
3. Click **Edit** to change the following parameters:
 - Type the IP address of the vCenter Server instance that contains the virtual machine data that you want to restore.
 - Type the user name and password that the Backup and Recovery capability will use to log in to vCenter Server.

Recovering single files from a virtual machine

After you have been given access to files on the destination virtual machine, you can use the Restore Agent window to view the backed up files created by Virtual Storage Console and find the ones that you need to restore from.

Before you begin

When you are restoring data, Restore Agent displays the first-available disk drive letter. If your guest operating system is Windows XP, 7, or 8, you must go to the Windows Disk Management snap-in application to manually assign the disk drive letter to other partitioned space on the disk. For more information about using the disk management utility, select **Help** from Disk Management.

Procedure

1. From the Restore Agent window, select a file from the **Disk tab**. Alternatively, click the **Backup** tab to find the Snapshot copy by name.
2. Right-click the name of the backup copy and select **Mount**. The **Backup** tab shows the drive letter for each backup copy only if the source virtual machine and destination virtual machine are the same. The contents of the backup copy are written to the new location.

What to do next

After finishing this task, you should clear the configuration cache.

Clear the configuration

After restoring your data from a backed up virtual machine, you should clear the configuration so you can upload another one later.

Procedure

1. From the Action pane in the Restore Agent window, click **Clear Configuration**.
2. In the resulting window, click **OK**. Backup and restore metadata will be removed from Restore Agent and Virtual Storage Console.

VSC CLI commands

The VSC for VMware vSphere command-line interface (CLI), which is labeled "VSC CLI" on your Windows desktop, provides you the benefits of a command-based view of the user interface. You can use this CLI to perform specific Backup and Recovery tasks, such as creating or deleting a backup of a virtual machine or datastore, as well as mounting a backup.

You should keep in mind the following information about the commands that you see in the interface:

- Virtual Storage Console commands are case-sensitive.
- There are no privilege levels; any user with a valid user name and password can run all commands.

For some commands, the following two parameters control the amount of output displayed:

verbose

This optional parameter provides detailed output when displaying information.

quiet

This optional parameter stops any output from displaying.

Note: Even with the quiet parameter specified, failed commands still display their failure messages.

Launching the VSC CLI

You can use either of two methods to launch the VS for VMware vSphere command-line interface (CLI), which is labeled "VSC CLI" on your Windows desktop. The first time you launch the VSC CLI, the application uses your Windows user credentials to grant you server access. Subsequent launches use stored credentials, speeding your access to the server.

About this task

When you issue your first CLI command, the CLI prompts you for your password and then runs the command. If the command succeeds, the CLI caches your user credentials and stores the information locally in an encrypted format.

Procedure

Double-click the **VSC CLI** icon or navigate to **Start > All Programs > IBM > IBM N series VSC CLI**.

An alternative method to using your Windows user credentials is to use the **smvi servercredential set** command to create custom user credentials that allow you to log in to the server.

smvi backup create

The **smvi backup create** command creates a backup of a virtual machine or datastore. You can also perform this operation using the VSC for VMware vSphere user interface.

Syntax

smvi backup create

```
-id {name | id} [name | id ...] [-backup-name {backup name}] [-server {server name}]  
[-include-independent] -exclude-datastores {name | id} [name | id ...] [-scripts  
{script name}] [-no-vmware-snapshot] [-update-mirror] [-quiet] [-verbose] [-user]  
[-help]
```

Parameters

[-id {name | id}] [name | id ...]

This mandatory parameter specifies the name or identification of the datastore or virtual machine that you are backing up. You can specify names or identifications of multiple datastores or virtual machines.

[-backup-name {backup name}]

This optional parameter specifies a backup copy name. After adding the flag, add a name for the backup copy. If you specify no name with this flag, the command fails. If you specify a name that contains only spaces, VSC for VMware vSphere automatically generates a name. If you specify a name that contains both spaces and other characters, VSC for VMware vSphere removes all leading and trailing spaces from the name.

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-include-independent]

This optional parameter specifies that datastores have only independent disks for a virtual machine are included in the backup.

[-exclude-datastores {name | id}] [name | id ...]

This optional parameter specifies the name or identification of the datastores or virtual machines to be excluded from the backup.

[-scripts {script name}]

This optional parameter specifies the name of the scripts to run with this backup.

[-no-vmware-snapshot]

This optional parameter prevents the creation of VMware snapshots of virtual machines during a backup.

[-update-mirror]

This optional parameter initiates a SnapMirror image on the secondary storage.

[-quiet]

This optional parameter stops any output from displaying.

[-verbose]

This optional parameter provides detailed output when displaying information.

`[-user]`

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

`[-help]`

This optional parameter displays help for this command.

Example: Creating a backup copy from a virtual machine

The following example creates a backup copy from a virtual machine named `nfs1_vm1` without specifying a backup name:

```
smvi backup create -id nfs1_vm1

[13:14] Starting backup request
[13:14] [WARN] Backup name is not set. Using default value 'backup_01fb4992a28188686d4e4a3ded34bfa4'
[13:14] Backing up datastore(s) ([nfs.123.17.170./vol/nfs_vol1/])
[13:14] Backing up the following virtual machine(s) ([nfs1_vm1])
[13:14] Creating VMware snapshots for all virtual machines that are being backed up.
[13:14] Creating storage snapshots for all datastores/virtual machines that are being backed up.
[13:14] Removing VMware snapshots for all virtual machines that are being backed up.
[13:14] Backup of datastores/virtual machines is complete.
SMVICLI-0100: Command completed successfully
```

smvi backup delete

The **smvi backup delete** command removes a virtual machine or datastore backup copy. You can also perform this operation using the VSC for VMware vSphere user interface.

Syntax

```
smvi backup delete -backup-name {backup name} [-server {server name}] [-quiet]
[-verbose] [-noprompt] [-user] [-help]
```

Description

When you delete the most recent backup associated with a backup job, then the Last Run Status value displayed for that backup job in the Schedule Backup Jobs window is that of the most recent remaining undeleted backup copy associated with the backup job.

Parameters

`[-backup-name {backup name}]`

This mandatory parameter specifies the backup copy you want to delete. After adding the flag, add the name of the backup copy.

`[-server {server name}]`

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is `localhost`.

`[-quiet]`

This optional parameter stops any output from displaying.

`[-verbose]`

This optional parameter provides detailed output when displaying information.

[-noprompt]

This optional parameter disables the default prompt that asks for confirmation when deleting a backup.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-help]

This optional parameter displays help for this command.

Example: Deleting a backup copy

The following example deletes a backup copy named new-one:

```
smvi backup delete -backup-name new-one
Are you sure you want to proceed and remove backup named 'new-one'?
[yes|NO] y
[15:15] Removed backup with name "new-one"
SMVICLI-0100: Command completed successfully
```

smvi backup list

The **smvi backup list** command displays information, such as the file path on a storage system to the Snapshot copy, about all of the created and saved backups within a virtual machine or datastore. You can also perform this operation using the VSC for VMware vSphere user interface.

Syntax

```
smvi backup list [-id {name | id} [name | id ...]] [-mounted] [-failed] [-recent]
[-with-vmware-snapshot] [-sfr-mounted] [-server {server name}] [-user] [-help]
```

Parameters

[-id {name | id} [name | id ...]]

This mandatory parameter specifies the name or identification of the datastores or virtual machines that you want to list.

[-mounted]

This optional parameter lists all mounted backups.

[-failed]

This optional parameter lists all failed backups. The default list is only successful backups.

[-recent]

This optional parameter lists the most recent backup.

[-with-vmware-snapshot]

This optional parameter lists the backups that were taken with a VMware snapshot.

[-sfr-mounted]

This optional parameter lists the backups that were mounted for SFR.

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-help]

This optional parameter displays help for this command.

Example: Listing backups in a datastore

The following example lists all of the backups within a datastore named data-store1:

```
smvi backup list -id data-store01

Id Name Date Entities Mounted VMware Snapshot Snapshot Name
-----
backup_sch_1_20090122233100 Jan 22, 2009 23:31 vmfs_vm1 No 10.72.248.38:/vol/kas1_102_
iscsi:smvi_backup_sch_1
20090122233100_36d2d99a-9ee0-4841-80c0-846698463e78_kas_sw_iscsi_ds1
```

smvi backup mount

The **smvi backup mount** command mounts a backup to verify its contents.

Syntax

smvi backup mount

```
-backup-name {backup name} -esx-server {esx server name}
[-server {server name}] [-quiet] [-verbose] [-user] [-help]
```

Privilege level

Note: To mount a VMFS datastore backup, the supplied ESX server must have SAN or iSAN access to the storage system, including required FC zoning or iSCSI discovery. To mount an NFS datastore backup, the supplied ESX server must be in the NFS export list of the original datastore.

Parameters

[-backup-name {backup name}]

This mandatory parameter specifies the backup you want to mount. After adding the flag, add the name of the backup.

[-esx-server {esx server name | IP address}]

This mandatory parameter specifies the name or IP address of the ESX server. This information describes where the backup resides on an ESX server.

Note: The server name is the name of the ESX server as viewed through the vSphere Client. This name might differ from the ESX server's host name or IP address.

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-quiet]

This optional parameter stops any output from displaying.

[-verbose]

This optional parameter provides detailed output when displaying information.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-help]

This optional parameter displays help for this command.

Example: Mounting a backup

The following example mounts a backup named `vmfs2_vm1` on an ESX server with the IP address of `123.12.1.23`:

```
smvi backup mount -backup-name vmfs2_vm1 -esx-server 123.12.1.23
```

```
[12:12] Starting mount request  
SMVICLI-0100: Command completed successfully
```

smvi backup rename

The **smvi backup rename** command changes the name of a backup. Changing the name of a backup also changes the name on the corresponding storage Snapshot copy on the associated IBM storage system. You can also perform this operation using the VSC for VMware vSphere user interface.

Syntax

smvi backup rename

-backup-name {backup name} -new-backup-name {new name} [-server {server name}]

[-user] [-help]

Parameters

[-backup-name {backup name}]

This mandatory parameter specifies the backup you want to rename. After adding the flag, add the name of the backup.

[-new-backup-name {new name}]

This mandatory parameter specifies the new name of the backup. After adding the flag, add a new name for the backup.

[-server-name {server name}]

This optional parameter specifies the name of the server to send the command to.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-help]

This optional parameter displays help for this command.

Example: Renaming a backup

The following example renames a backup named `vmfs2-vm1` to `volume-2`:

```
smvi backup rename -backup-name vmfs2-vm1 -new-backup-name volume-2
```

```
[15:52] Backup "vmfs2-vm1" has been renamed to "volume-2"  
SMVICLI-0100: Command completed successfully
```

smvi backup restore

The **smvi backup restore** command enables you to restore a virtual machine or datastore from a backup copy. You can also perform this operation using the VSC for VMware vSphere user interface.

Syntax

smvi backup restore

```
-id {name | id} [-esx-server {esx server name}] [-backup-name {backup name}]  
[-vmdk {hard disk name}] [-server {server name}] [-restart-vm] [-quiet]  
[-verbose] [-noprompt] [-user] [-help]
```

Parameters

[-id {name | id}]

This mandatory parameter specifies the name or identification of the datastore or virtual machine that you are restoring.

[-esx-server {esx server name | IP address}]

This mandatory parameter specifies the name or IP address of the ESX server. The parameter is required when restoring a VMFS datastore, or a virtual machine that resides on a VMFS datastore, as well as when restoring an NFS virtual machine. The server name is the name of the ESX server as viewed through the vSphere Client. This name might differ from the host name or IP address of the ESX server.

[-backup-name {backup name}]

This optional parameter specifies which backup to restore. After adding the flag, you can add the name of the backup. If not specified, the latest available backup for the specified datastore or virtual machine is restored.

[-vmdk {hard disk name}]

This optional parameter specifies which hard disks are to be restored.

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-restart-vm]

This optional parameter restarts the virtual machine after the restore operation.

[-quiet]

This optional parameter stops any output from displaying.

[-verbose]

This optional parameter provides detailed output when displaying information.

[-noprompt]

By default, a prompt appears, asking for confirmation when restoring a backup. This optional parameter disables the prompt.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-help]

This optional parameter displays help for this command.

Example: Restoring a backup

The following example restores a virtual machine named `nfs1-vm1` from a backup named `backup-411`:

```
smvi backup restore -id nfs1_vm1 -backup-name backup-411

Are you sure you want to proceed with this operation? [yes|NO] y
[11:04] Starting restore request
[11:04] [WARN] No active mounts found for datastore vmfs_ds1
(47ab69d8-e7c72da0-d6c5-001a6412251d)
[11:05] Restoring nfs virtual machine on folder 'nfs1_vm1'
[11:07] Reloading virtual machine
[11:07] Restore is complete
SMVICLI-0100: Command completed successfully
```

smvi backup unmount

The **smvi backup unmount** command unmounts a mounted virtual machine or datastore backup.

Syntax

```
smvi backup unmount
-backup-name {backup name} [-server {server name}] [-quiet] [-verbose]
[-user] [-help]
```

Description

Note: You must unmount a mounted backup in order to delete the backup or any of its preceding Snapshot copies.

Parameters

[-*backup-name* {*backup name*}]

This mandatory parameter specifies which backup to unmount. After adding the flag, add the name of the backup.

[-*server* {*server name*}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-*quiet*]

This optional parameter stops any output from displaying.

[-*verbose*]

This optional parameter provides detailed output when displaying information.

[-*user*]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-*help*]

This optional parameter displays help for this command.

Example: Unmounting a backup

The following example unmounts a backup named `vmfs2_db`:

```
smvi backup unmount -backup-name vmfs-2-db  
[11:55] Starting unmount request  
[11:55] Unmount is complete  
SMVICLI-0100: Command completed successfully
```

smvi discover datastores

The **smvi discover datastores** command lists the datastores that are managed by the current vCenter Server and that reside on the storage systems currently assigned to your SnapManager for Virtual Infrastructure server.

Syntax

```
smvi discover datastores [-help]
```

Parameters

[-help]

This optional parameter displays help for this command.

Example: Listing the datastores

The following example lists all the datastores managed by the current vCenter Server that reside on storage systems assigned to SnapManager for Virtual Infrastructure:

```

smvi discover datastores

Password for Nseries\vanib: *****
Datacenter: Aladdin
Datastore: nfs_datastore6
NFS: 172.17.170.21:/vol/nfs_vol6
Datastore: nfs_datastore7
NFS: 172.17.170.21:/vol/nfs_vol7
Datastore: nfs_datastore7 (Backup test1)
NFS: 172.17.170.21:/vol/nfs_vol7_mount_33e49878c5e74363825e84652a724aef
Datastore: nfs_datastore7 (Backup test0)
NFS: 172.17.170.21:/vol/nfs_vol7_mount_90a6b1e7d6f948beaa6735af9692b3d4
Datastore: nfs_datastore7 (Backup backup_fgfdgfdgf_20080707134801)
NFS: 172.17.170.21:/vol/nfs_vol7_mount_e50fc0eda0674cfbbf200f87f83ba8eb
Datastore: nfs_datastore8
NFS: 172.17.170.21:/vol/nfs_vol8
Datastore: nfs_datastore8 (Backup
backup_7d8597b0dffffd5c81806728dd45aea48)
NFS: 172.17.170.21:/vol/nfs_vol8_mount_e7df47fbde00446cb6b589c821adc4dd
Datastore: vmfs_datastore5
LUN: 172.17.170.21:/vol/vmfs_vol5/lun5 Partition: 1
LUN: 172.17.170.21:/vol/vmfs_vol6/lun6 Partition: 1
Datastore: vmfs_datastore2
LUN: 172.17.170.21:/vol/vmfs_vol2/vmfs_lun2 Partition: 1
Datastore: vmfs_datastore3
LUN: 172.17.170.21:/vol/vmfs_vol3/lun3 Partition: 1
Datastore: vmfs_datastore4
LUN: 172.17.170.21:/vol/vmfs_vol4/lun4 Partition: 1
Datastore: vmfs_datastore7
LUN: 172.17.170.21:/vol/vmfs_vol7/qtrees_vol7/lun7 Partition: 1
Datastore: snap-00000002-vmfs_datastore
LUN: 172.17.170.21:/vol/vmfs_vol1/vmfs_lun1 Partition: 1
Datastore: vmfs7_testAJ_1
LUN: 172.17.170.21:/vol/volaj1/lun1 Partition: 1
Datastore: vmfs7_testAJ-2
LUN: 172.17.170.21:/vol/volaj1/lun2 Partition: 1
Datacenter: Bellagio

```

smvi filerestore add-portgroup

The **smvi filerestore add-portgroup** command assigns virtual machines to a port group. You can also perform this operation using the VSC for VMware vSphere user interface.

Syntax

```

smvi filerestore add-portgroup
[-name{port group name}] [-server{server name}] [-user]
[-verbose] [-help]

```

Parameters

[-name {port group name}]

This mandatory parameter specifies the name of the port group, or network, that is used to enable or disable administrator-assisted file-level restore operations.

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-verbose]

This optional parameter provides detailed output when displaying information.

[-help]

This optional parameter displays help for this command.

smvi filerestore delete-portgroup

The **smvi filerestore delete-portgroup** command removes the port group and disables file restore sessions for the virtual machines assigned to the port group. You can also perform this operation using the VSC for VMware vSphere user interface.

Syntax

smvi filerestore delete-portgroup

[-name {port group name}] [-server {server name}] [-user]
[-verbose] [-help]

Parameters

[-name {port group name}]

This mandatory parameter specifies the name of the port group, or network, that is used to enable or disable administrator-assisted file-level restore operations.

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-verbose]

This optional parameter provides detailed output when displaying information.

[-help]

This optional parameter displays help for this command.

smvi notification list

The **smvi notification list** command displays information about the alert notification.

Syntax

smvi notification list *[-server {server name}] [-user] [-verbose] [-help]*

Parameters

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-verbose]

This optional parameter provides detailed output when displaying information.

[-help]

This optional parameter displays help for this command.

smvi notification set

The **smvi notification set** command displays information about the alert notification.

Syntax

smvi notification set

[-smtp server {dns name | ip address}] [-from {from email address}] [-to {to email address}] [-server {server name}] [-user] [-verbose] [-help]

Parameters

[-smtp server {dns name | ip address}]

This mandatory parameter specifies the name or IP address of the SMTP server that handles the test notification e-mail.

[from {from email address}]

This mandatory parameter specifies the sender e-mail address.

[to {to email address}]

This mandatory parameter specifies the comma-separated list of recipient e-mail addresses.

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-verbose]

This optional parameter provides detailed output when displaying information.

[-help]

This optional parameter displays help for this command.

smvi notification test

The **smvi notification test** command displays information about the test notification.

Syntax

smvi notification test ***[-server {server name}] [-user] [-verbose] [-help]***

Parameters

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-verbose]

This optional parameter provides detailed output when displaying information.

[-help]

This optional parameter displays help for this command.

smvi restoreagent set

The **smvi restoreagent set** command sets the default installation URL of the restore agent. You can also perform this operation using the VSC for VMware vSphere user interface.

Syntax

```
smvi restoreagent set [-url] [-server {server name}] [-user] [-verbose] [-help]
```

Parameters

[-url]

This mandatory parameter provides an URL that points to a customer location for the restore agent installer.

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-verbose]

This optional parameter provides detailed output when displaying information.

[-help]

This optional parameter displays help for this command.

smvi servercredential delete

The **smvi servercredential delete** command deletes a user account created by the **smvi servercredential set** command.

Syntax

```
smvi servercredential delete -username {user name} [-help]
```

Description

Note: You cannot run this command from a remote host.

Parameters

[-username {user name}]

This mandatory parameter specifies the internal user account that you want to delete.

[-help]

This optional parameter displays help for this command.

Example: Deleting a user account

The following example deletes the olduser2 user account:

```
smvi servercredential delete -username olduser2
SMVICLI-0100: Command completed successfully
```

smvi servercredential list

The **smvi servercredential list** command lists a user account created by the **smvi servercredential set** command.

Syntax

```
smvi servercredential list [-help]
```

Description

Note: You cannot run this command from a remote host.

Parameters

[-help]

This optional parameter displays help for this command.

Example: Listing the server credentials

The following example lists the current SnapManager for VI server credentials:

```
smvi servercredential list
Username
-----
administrator
```

smvi servercredential set

The **smvi servercredential set** command adds a user account for Backup and Recovery capability to use for authentication instead of your Windows user credentials.

Syntax

```
smvi servercredential set [-help]
```

Description

Note: You cannot run this command from a remote host.

Parameters

[-help]

This optional parameter displays help for this command.

Example: Adding a user account

The following example adds a user account named administrator and sets a seven character password:

```
smvi servercredential set
Username: administrator
Password: *****
SMVICLI-0100: Command completed successfully
```

smvi storagesystem add

The **smvi storagesystem add** command adds an IBM storage system to your configuration. You can also perform this operation using the VSC for VMware vSphere user interface.

Syntax

```
smvi storagesystem add
  -name {DNS name | IP address}    [-server {server name}]
  [-user] [-help]
```

Parameters

[-name {DNS name | IP address}]

This mandatory parameter specifies the DNS name or management IP address of the IBM storage system that you are adding.

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-help]

This optional parameter displays help for this command.

Example: Adding a storage system

The following example adds a server with an IP address of 123.18.20.25, enters its administrator user name as client-1, and enters this user's 10-character password:

```
smvi storagesystem add -name 123.18.20.25
Enter username : client-1
Enter password : *****
smvicli-0100: Command completed successfully
```

smvi storagesystem delete

The **smvi storagesystem delete** command deletes an IBM storage system. You can also perform this operation using the VSC for VMware vSphere user interface.

Syntax

```
smvi storagesystem delete  
-name {DNS name | IP address} [-server {server name}]  
[-user] [-help]
```

Parameters

[-name {DNS name | IP address}]

This mandatory parameter specifies the DNS name or management IP address of the IBM storage system that you are deleting. You must provide the exact name or IP address of the storage system, or the command fails.

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-help]

This optional parameter displays help for this command.

Example: Deleting a storage system

The following example deletes a IBM storage system named Jaguar:

```
smvi storagesystem delete -name Jaguar  
smvicli-0100: Command completed successfully
```

smvi storagesystem list

The **smvi storagesystem list** command lists the added IBM storage systems.

Syntax

```
smvi storagesystem list  
[-server {server name}] [-user] [-verbose] [-help]
```

Parameters

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-verbose]

This optional parameter provides detailed output when displaying information.

[-help]

This optional parameter displays help for this command.

Example: Listing the storage systems

The following example lists the IBM storage systems that reside in the local SnapManager for VI server; in this case, a single storage system with an IP address of 123.17.170.21:

```
smvi storagesystem list
```

Name	IP Address
-----	-----
123.17.170.21	123.17.170.21

smvi storagesystem modify

The **smvi storagesystem modify** command modifies a saved IBM storage system. You can also perform this operation using the VSC for VMware vSphere user interface.

Syntax

```
smvi storagesystem modify -name {DNS name | IP address}  
[-server {server name}] [-user] [-help]
```

Parameters

[-name {DNS name | IP address}]

This mandatory parameter specifies the DNS name or management IP address of the IBM storage system that you are modifying. You must provide the exact name or IP address of the storage system, or the command fails.

[-server {server name}]

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

[-user]

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

[-help]

This optional parameter displays help for this command.

Example: Modifying a storage system

The following example updates the user name for the IBM storage system with an IP address of 123.18.20.25 to root and enters this user's ten character password:

```
smvi storagesystem modify -name 123.18.20.25
```

```
Username for 123.18.20.25: root
```

```
Password for 123.18.20.25: *****
```

```
smvicli-0100: Command completed successfully
```

smvi version

The **smvi version** command displays the version of the VSC for VMware vSphere CLI and the SnapManager for VI server.

Syntax

```
smvi version [-server {server name}] [-help]
```

Parameters

`[-server {server name}]`

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

`[-help]`

This optional parameter displays help for this command.

Example: Displaying the version

The following example displays the VSC for VMware vSphere CLI and server version:

```
smvi version
SnapManager for Virtual Infrastructure CLI Rballys.4N_120127_0000 (Build: 120127
)
SnapManager for Virtual Infrastructure Server Rballys.4N_120127_0000 (Build: 120
127)
```

Programmable APIs

VSC for VMware vSphere supports both a VSC for VMware vSphere Application Programming Interface (API) and a Provisioning and Cloning API.

What the programmable APIs are

Virtual Storage Console for VMware vSphere provides programmable application interfaces (APIs) for Provisioning and Cloning and for VMware vCloud. The APIs are exposed using Simple Object Access Protocol (SOAP). They provide a layer above the NetApp Manageability SDK, the VMware VI SDK, and the VMware vCloud SDK, but do not require any of these in the consumer application or script.

What you can do with the APIs for VMware vCloud

The Virtual Storage Console for VMware vSphere provides APIs that enable you to manage credentials for multiple vCenter Servers, discover vCloud Director objects for vCloud tenants, and provision and clone vApps, provided that you have appropriate vCloud Director privileges.

Provisioning and Cloning programmable API

The Provisioning and Cloning Application Programming Interface (API) is designed to be leveraged with the VI SDK. It provides end-to-end automated datastore provisioning and offloads the intricacies of storage object cloning while cloning virtual machines.

The managed object reference returned by the VMware VI SDK is used to identify components in the vCenter Inventory. You can view this information using the Managed Object Browser on the vCenter Server.

This version of the Provisioning and Cloning API exposes the virtual machine clone creation engine (which includes the redeploy feature), the datastore management engine (create, destroy, resize), and the file copy/clone offload engine. There are also two general-purpose utility methods included:

- `getVmFiles` returns a list of files that make up the virtual machine. This is useful for creating the list of files required in the `cloneSpec` API.
- `getMoref` returns the managed object reference of the requested object based on name and type. The `getMoref` returns the first object that matches the name and type combination. For this reason, this method should not be used in production environments unless all object names are unique.

The virtual machine clone engine

The virtual machine clone engine provides two clone creation and routing methods: `createClones` and `redeployVMs`.

- `createClones` can be used to create virtual machine clones on new or existing datastores. When more than one datastore is created, the FlexClone feature on the controller is leveraged to create clones of the datastore.
- `redeployVMs` provides the ability to redeploy the virtual hard drives of the source virtual machine to the virtual machines specified. This feature leverages the FlexClone feature on the controller as well.

The datastore management engine

The datastore management engine provides three methods for managing datastores: createDatastore, resizeDatastore, and destroyDatastore.

- The createDatastore method provides the ability to provision storage on the controller, present it to the ESX hosts, and create a datastore.
- The resizeDatastore method provides the ability to grow and shrink NFS-based datastores and grow VMFS-based datastores.
- The destroyDatastore method provides the ability to delete all virtual machines associated with the datastore, unmount it from ESX hosts, destroy the storage objects on the controller, and free the space.

The file copy/clone offload engine

The file copy/clone offload engine provides four methods. These methods provide the ability to execute and monitor file copy and clone operations.

This engine provides the ability to offload file copy and clone operations to the controller for NFS-based datastores. This functionality is unique compared to that provided by the other engines in that it does not require a Virtual Center session. An ESX host session can be used instead.

The input to the methods is a combination of complex (specification and message) and simple (string, int, long, boolean, and so on) data types. The specifications and messages are described below.

Note: Very little verification or validation is done in the API. For example, if there is not enough space to create the requested datastore(s), the API method will fail.

Provisioning and Cloning methods

This section describes all the available Provisioning and Cloning methods.

Virtual machine clone creation and redeploy engine

This section describes the APIs for interfacing with the virtual machine clone creation and redeploy engine.

createClones:

You can use the createClones method to create virtual machine clones on new or existing datastores.

The source can be a virtual machine or a virtual machine template. The source can be further refined by specifying a virtual machine snapshot. The following options cause at least one native clone (built into Virtual Center) operation to occur:

- clone source is powered on
- virtual machine snapshot is specified
- hard drive transformation is specified

The virtual machine or template must not contain any RDMs, must not contain any devices that use VMDirectPath, and must be connected.

The mix of VirtualIDEController attached hard drives and VirtualSCSIController hard drives in the same virtual machine may result in the drives being reordered in the resulting clones, therefore this is not supported. The creation of virtual machines based on hardware version vmx-07 will fail on ESX 3.5 hosts.

Status

Current (added in version 2.1)

Type

Asynchronous

Parameters

Name	Type	Value	Description
requestSpec	Object	RequestSpec	Request can specify a vCenter server only. This method does not support direct connections to ESX hosts. Note: See RequestSpec.

XML

```
complexType name="createClones">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://server.kamino.ibm.com/}requestSpec"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Returns

Type	Value	Description
String	Task:task-2	A managed object reference to a vCenter task. This task can be monitored and altered using the VI SDK

Return XML

```
<complexType name="createClonesResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

RedeployVMs:

You can redeploy the virtual hard drives associated with a virtual machine to other virtual machines.

The source can be a virtual machine, a virtual machine template, or a virtual machine snapshot. The following options cause a native (built into Virtual Center) clone operation before it can use the rapid clone methodology:

- clone source is powered on
- virtual machine snapshot is specified
- hard drive transformation is specified

The virtual machine or template must not contain any RDMs or any devices that use VMdirectPath, and must be in a good state (connected).

Status

Current (added in version 3.0)

Type

Asynchronous

Parameters

Name	Type	Value	Description
requestSpec	Object	RequestSpec	Note: See RequestSpec.

XML

```
complexType name="redeployVMs">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://server.kamino.ibm.com/}requestSpec"
minOccurs="0"/>
        <element name="arg1" type="{http://server.kamino.ibm.com/}controllerSpec"
maxOccurs="unbounded" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Returns

Type	Value	Description
String	Task:task-2	A managed object reference to a vCenter task. This task can be monitored and altered using the VI SDK

Return XML

```
<complexType name="redeployVMsResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/2001/XMLSchema}string"
minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Datastore management engine

This section describes the APIs for interfacing with the datastore management engine.

createDatastore:

You can use the createDatastore method to provision storage on the storage controller, attach it to one or more ESX hosts and create a datastore.

More than one ESX host can be chosen by specifying the managed object reference of a cluster or datacenter in the DatastoreSpec.

Status

Current (added in version 3.0)

Type

Synchronous

Parameters

Name	Type	Value	Description
dsSpec	Object	DatastoreSpec	The specification describing the datastore to create. Note: See DatastoreSpec.
requestSpec	Object	RequestSpec	Request can specify a vCenter server only. This method does not support direct connections to ESX hosts. Note: See RequestSpec.

XML

```
complexType name="createDatastore">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://server.kamino.ibm.com/}requestSpec"
          minOccurs="0"/>
        <element name="arg1" type="{http://server.kamino.ibm.com/}controllerSpec"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Returns

Type	Value	Description	
String	newDatastore	The name of the new datastore that was created.	

Return XML

```
<complexType name="createDatastoreResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

resizeDatastore:

You can use the `resizeDatastore` method to grow or shrink NFS-based datastores (and associated storage objects on the controller), and grow VMFS-based datastores (and associated storage objects on the controller).

Status

Current (added in version 3.0)

Type

Synchronous

Parameters

Name	Type	Value	Description
dsSpec	Object	DatastoreSpec	Specification describing datastore resize request.
requestSpec	Object	RequestSpec	Request can specify a vCenter server only. This method does not support direct connections to ESX hosts.

XML

```
complexType name="resizeDatastore">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://server.kamino.ibm.com/}requestSpec"
          minOccurs="0"/>
        <element name="arg1" type="{http://server.kamino.ibm.com/}datastoreSpec"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Returns

Type	Value	Description
String	Task:task-2	A managed object reference to a vCenter task.

Return XML

```
complexType name="resizeDatastoreResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

destroyDatastore:

You can use the `destroyDatastore` method to delete any virtual machine with a file on the specified datastore, delete the datastore (after detaching from each ESX host), take the storage objects offline, and destroy the datastore (to free space).

Status

Current (added in version 3.0)

Type

Synchronous

Parameters

Name	Type	Value	Description
dsSpec	Object	DatastoreSpec	Specification describing datastore resize request. Note: See DatastoreSpec.
requestSpec	Object	RequestSpec	Request can specify a vCenter server only. This method does not support direct connections to ESX hosts. Note: See RequestSpec.

XML

```
complexType name="destroyDatastore">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://server.kamino.ibm.com/}datastoreSpec"
          minOccurs="0"/>
        <element name="arg1" type="{http://server.kamino.ibm.com/}requestSpec"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

```

    minOccurs="0"/>
  </sequence>
</restriction>
</complexContent>
</complexType>

```

Returns

Type	Value	Description
String	Task:task-2	A managed object reference to a vCenter task.

Return XML

```

<complexType name="destroyDatastoreResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

Connection Broker features

This section describes the APIs for interfacing with connection broker features.

performViewImport:

You can use the performViewImport method to import the specified virtual machines into a VMware View Server.

Status

Current (added in version 3.2)

Type

Synchronous

Parameters

Name	Type	Value	Description
dsSpec	Object	DatastoreSpec	The specification describing the connection broker information.
requestSpec	Object	RequestSpec	Request can specify a vCenter server only. This method does not support direct connections to ESX hosts.

Name	Type	Value	Description
vmsForImport	List <String>		A list of the virtual machines (by name) that should be imported into the View server.

XML

```
complexType name="createDatastore">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://server.kamino.ibm.com/}connectionBrokerSpec"
          minOccurs="0"/>
        <element name="arg1" type="{http://server.kamino.ibm.com/}requestSpec"
          minOccurs="0"/>
        <element name="arg2" type="{http://server.kamino.ibm.com/}vmsForImport"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Returns

Type	Value	Description
Void	N/A	Nothing returned

Return XML

```
complexType name="createDatastore">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/2001/XMLSchema}void"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

createXenImportFile:

You can use the createXenImportFile method to create a file for importing virtual machines into a Citrix XenDesktop server.

Status

Current (added in version 3.2)

Type

Synchronous

Parameters

Name	Type	Value	Description
dsSpec	Object	ConnectionBrokerSpec	The specification describing the connection broker information.
requestSpec	Object	RequestSpec	Request can specify a vCenter server only. This method does not support direct connections to ESX hosts.
vmsForImport	List		A list of the virtual machines (by name) that should be imported into the View server.

XML

```
complexType name="createXenImportFile">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://server.kamino.ibm.com/}connectionBrokerSpec"
          minOccurs="0"/>
        <element name="arg1" type="{http://server.kamino.ibm.com/}requestSpec"
          minOccurs="0"/>
        <element name="arg2" type="{http://server.kamino.ibm.com/}vmsForImport"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Returns

Type	Value	Description
Void	N/A	Nothing returned

Return XML

```
complexType name="createXenImportFileResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/2001/XMLSchema}void"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Copy/Clone offload engine

This section describes the APIs for interfacing with the Copy/Clone offload engine.

The fileCopyOffload and fileCloneOffload take **VmFileSpec** and **RequestSpec** as arguments. The DatastoreSpec should contain only the datastore-managed object reference and a reference to the controller. This is defined as *Existing Datastore* in the formulas section of the *DatastoreSpec* documentation.

fileCopyOffload:

You can use the `fileCopyOffload` method to offload the copy of an NFS datastore file to the controller. This method should be used in cases where a full copy (all unique blocks) is required. In all other cases, the `fileCloneOffload` should be used.

This process involves a start-up time, which is quickly recovered when copying large files (because the offloaded controller base copy is very efficient). This start-up time may cause the offloaded copy of small files to take longer than using a host-based copy.

This method supports copying a file within the same controller. The `VmFileSpec` for the source and destination must specify the same controller.

Status

Current (added in version 3.0)

Type

Asynchronous

Parameters

Name	Type	Value	Description
source	Object	VmFileSpec	Specification describing the source file (datastore and controller).
destination	Object	VmFileSpec	Specification describing the destination file (datastore and controller).
requestSpec	Object	RequestSpec	Request can specify a vCenter server or ESX host.

XML

```
complexType name="fileCopyOffload">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://server.kamino.ibm.com}/vmFileSpec"
          minOccurs="0"/>
        <element name="arg1" type="{http://server.kamino.ibm.com}/vmFileSpec"
          minOccurs="0"/>
        <element name="arg2" type="{http://server.kamino.ibm.com}/requestSpec"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Returns

Type	Value	Description
Integer	876234	The operation identifier to monitor using getFileOpOffloadStatus.

Return XML

```
complexType name="fileCopyOffloadResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/2001/XMLSchema}int"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

fileCloneOffload:

You can use the fileCloneOffload method to offload the clone of an NFS datastore file to the controller.

This process uses the file level FlexClone feature of the controller. This process automatically falls back to the controller offloaded copy as needed if the **failBackToCopy** parameter is set to true.

This method supports only cloning the file within the same volume. If **failBackToCopy** is set to true, this method supports copying file within the same controller. In both cases, the VmFileSpec for the source and destination must specify the same controller.

The most effective use of this method is to employ a strategy where the output of the first operation (the destination file) becomes the input (the source file) for the next operation. For example, to create three clones of test-flat.vmdk, the following process (pseudo code) is the most efficient:

```
clone(test-flat.vmdk, test1-flat.vmdk)
clone(test1-flat.vmdk, test2-flat.vmdk)
clone(test2-flat.vmdk, test3-flat.vmdk)
```

Status

Current (added in version 3.0)

Type

Asynchronous

Parameters

Name	Type	Value	Description
source	Object	VmFileSpec	Specification describing the source file (datastore and controller).

Name	Type	Value	Description
destination	Object	VmFileSpec	Specification describing the destination file (datastore and controller).
fallBackToCopy	Boolean		If set to true, engine runs in "fully automatic mode" which falls back to an offloaded copy as needed. If false, conditions that would normally fall back to a copy result in an error (which the caller must deal with).
requestSpec	Object	RequestSpec	Request can specify a vCenter server or ESX host.

XML

```

complexType name="fileCloneOffload">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://server.kamino.ibm.com/}vmFileSpec"
          minOccurs="0"/>
        <element name="arg1" type="{http://server.kamino.ibm.com/}vmFileSpec"
          minOccurs="0"/>
        <element name="arg2" type="{http://server.kamino.ibm.com/}requestSpec"
          minOccurs="0"/>
        <element name="arg3" type="{http://www.w3.org/2001/XMLSchema}boolean"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

Returns

Type	Value	Description
Integer	876234	The operation identifier to monitor using getFileOpOffloadStatus.

Return XML

```

complexType name="fileCloneOffloadResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/2001/XMLSchema}int"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

getFileOpOffloadStatus:

You can use the `getFileOpOffloadStatus` method to track the progress of a `fileCopyOffload` or `fileCloneOffload` operation.

The status will be `complete`, `failed`, or `running`. When this method returns a **StatusMessage** with a status of `complete` or `failed`, the operation information is marked for cleanup, which occurs five minutes later. After the operation information has been cleaned up, it is no longer visible using this method. The progress field displays information about the progress of the operation.

Status

Current (added in version 3.0)

Type

Synchronous

Parameters

Name	Type	Value	Description
<code>opId</code>	Integer	876234	The operation identifier returned from <code>fileCopyOffload</code> or <code>fileCloneOffload</code> .

XML

```
complexType name="getFileOpOffloadStatus">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://www.w3.org/2001/XMLSchema}int"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Returns

Type	Value	Description
String	StatusMessage	Information describing status, progress, and reason for error (if operation fails). Note: See StatusMessage.

Return XML

```
complexType name="getFileOpOffloadStatusResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://server.kamino.ibm.com/}statusMessage"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

clearAllFinishedOpOffloadStatus:

You can use the clearAllFinishedOpOffloadStatus method to start the cleanup timer described in getFileOpOffloadStatus for all operations that have a status of complete or failed.

Status

Current (added in version 3.0)

Type

Synchronous

Parameters

Name	Type	Value	Description
opId	Integer		The operation identifier returned from fileCopyOffload or fileCloneOffload.

XML

```
complexType name="clearAllFinishedOpOffloadStatus">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Returns

Void

Return XML

```
complexType name="clearAllFinishedOpOffloadStatusResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Utility methods

This section describes utility methods that return output, such as a list of created virtual machines and the managed object reference of each virtual machine.

getVms:

You can use the getVms method to obtain the list of virtual machines that were created using the createClones method. This list can be used in the redeployVMs method.

Status

Current (added in version 3.0)

Type

Synchronous

Parameters

Name	Type	Value	Description
vmMorRef	Object	The managed object reference of the VM.	The managed object reference of the VM.

XML

```
complexType name="getVms">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="arg1" type="{http://server.kamino.ibm.com/}requestSpec"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Returns

Type	Value	Description
vmMorRef	String	The managed object reference of the VM.
requestSpec	RequestSpec	Request can specify a vCenter server or ESX host.

Return XML

```
<complexType name="getVmsResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/2001/XMLSchema}string"
          maxOccurs="unbounded" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

getVmFiles:

You can use the getVmFiles method to obtain a skeleton list of VmFileSpec to be completed and used in the submission to createClones.

Status

Current (added in version 2.1)

Type

Synchronous

Parameters

Name	Type	Value	Description
vmMorRef	String		The managed object reference of the VM.
requestSpec	Object	RequestSpec	Request can specify a vCenter server or ESX host.

XML

```
complexType name="getVmFiles">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="arg1" type="{http://server.kamino.ibm.com/}requestSpec"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Returns

Type	Value	Description
List	Object <VmFileSpec>	A list of VmFileSpec based on the VM specified. This information should be modified and submitted using the CloneSpec.

Return XML

```
<complexType name="getVmFilesResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://server.kamino.ibm.com/}vmFileSpec"
          maxOccurs="unbounded" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

getMoref:

You can use the `getMoref` method to obtain the managed object reference of the requested object based on name and type.

The `getMoref` method returns the first object that matches the name and type combination. For this reason, this method should not be used in production environments unless all object names are unique.

Status

Current (added in version 2.1)

Type

Synchronous

Parameters

Name	Type	Value	Description
name	String		Name of object to look for.
type	String		Managed object type.
requestSpec	String		Request can specify a vCenter server or ESX host.

XML

```
complexType name="getMoref">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="arg1" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="arg2" type="{http://server.kamino.ibm.com/}requestSpec"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Returns

Type	Value	Description
String		Managed object reference in string format.

Return XML

```
<complexType name="getMorefResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Provisioning and Cloning specifications and messages

The Provisioning and Cloning API provides several specifications and messages.

RequestSpec

The RequestSpec specification describes the URL of the VMware vCenter SDK or ESX host as well as the authentication information. The authentication information may be in the form of a user name and password combination or a VMware Session ID. An optional clone specification may also be present.

Properties

Type	Value	Description
serviceUrl	String	URL for the VMware vCenter SDK
vcUser	String	VMware vCenter username (null ok if using vcSession)
vcPassword	String	VMware vCenter password (null ok if using vcSession)
vcSession	String	VMware session (null ok if using vcUser/vcPassword)
cloneSpec	Object CloneSpec	A clone specification

Notes

- **cloneSpec** may be null when using this spec with anything other than redeployVMs or createClones.
- **vcSession** should be null if **vcUser** and **vcPassword** are used.
- **vcUser** and **vcPassword** should be null if **vcSession** is used.

XML

```
complexType name="requestSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="cloneSpec" type="{http://server.ibm.com/}cloneSpec"
          minOccurs="0"/>
        <element name="serviceUrl" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="vcPassword" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="vcSession" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="vcUser" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

CloneSpec

The CloneSpec specification describes a request to create clones of a virtual machine or template or to redeploy the virtual hard drives.

When used with the redeployVMs method, each virtual machine in the map named clones will have its virtual hard drives replaced with those of the source. The string in this map is the name of the virtual machine to be redeployed and the VmSpec describes this virtual machine.

When CloneSpec is used with the createClones method, a new virtual machine is created for each entry in the clones map. The string in this map is the name of the

clone and the VmSpec describes the new clone configuration. The list named files describes the files that make up the source virtual machine or template. This list can be used to specify different destinations for each file as well as to create new datastores.

Type	Value	Description
templateMoref	String	Source VM or template of cloning operation. String representation of type and value of ManagedObjectReference from VMware VI API.
snapshotMoref	String	The managed object reference for a snapshot of the source virtual machine to base the clones on.
containerMoref	String	Destination for resulting clones. Valid destination types: Datacenter, ResourcePool, ClusterComputeResource, and ComputeResource, A string representation of type and value of ManagedObjectReference from the VMware VI API.
destVmFolderMoref	String	Virtual machine folder the clones should be created in. If null, clones are created at the root virtual machine folder.
vmTransform	String	Transforms all virtual hard drives to specified format. Should be specified only when there is actual work to do. Specifying a transform when one is not required causes unnecessary work. Options are null, flat, and sparse.
hardwareVersion	String	Upgrade hardware version from a previous version to vmx-04 or vmx-07. Note: vmx-04 is supported by ESX 3.5 and both are supported by ESX 4.0.
clones	Map <String, VmSpec>	Map of new virtual machine name to virtual machine specification (VmSpec).
files	List <VmFileSpec>	List of files that make up source virtual machine or template specified in templateMoref .

Type	Value	Description
memMB	Long	Override the source virtual machine (or template) amount of memory during cloning process. Value is in MB.
numberCPU	Int	Override the source virtual machine (or template) number of CPUs during cloning process.

XML

```

<complexType name="cloneSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="clones">
          <complexType>
            <complexContent>
              <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
                <sequence><element name="entry" maxOccurs="unbounded" minOccurs="0">
                  <complexType>
                    <complexContent>
                      <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
                        <sequence>
                          <element name="key" type="{http://www.w3.org/2001/XMLSchema}string"
                            minOccurs="0"/>
                          <element name="value" type="{http://server.ibm.com/}vmSpec"
                            minOccurs="0"/>
                        </sequence>
                      </restriction>
                    </complexContent>
                  </complexType>
                </element>
              </sequence>
            </restriction>
          </complexContent>
        </element>
        <element name="connBroker" type="{http://server.ibm.com/}connectionBrokerSpec"
          minOccurs="0"/>
        <element name="containerMoref" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="destVmFolderMoref" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="files" type="{http://server.kamino.ibm.com/}vmFileSpec"
          maxOccurs="unbounded" minOccurs="0"/>
        <element name="hardwareVersion" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="snapshotMoref" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="templateMoref" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="vmTransform" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="memMB" type="{http://www.w3.org/2001/XMLSchema}long"
          minOccurs="0"/>
        <element name="numberCPU" type="{http://www.w3.org/2001/XMLSchema}int"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

VmFileSpec

The VmFileSpec specification describes the source configuration file (vmx) or the source virtual hard disk files (vmdk) as well as the destination datastore specification.

Properties

Type	Value	Description
sourcePath	String	Path to vmx or vmdk file. The string Configuration File can be passed in place of an actual vmx file.
destDatastoreSpec	DatastoreSpec	Destination datastore specification.

XML

```
<complexType name="vmFileSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="destDatastoreSpec" type="{http://server.ibm.com/}datastoreSpec"
          minOccurs="0"/>
        <element name="sourcePath" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

DatastoreSpec

The DatastoreSpec specification describes the destination datastore. This can describe new datastores to be created as well as existing datastores.

See the "Required parameters" section below for valid parameter combinations. The number of clones must be evenly divisible by the number of datastores.

Properties

Type	Value	Description
mor	String	Destination datastore. String representation of type and value of ManagedObjectReference from the VMware VI API.
targetMor	String	The managed object reference of the vCenter object in which to attach the new datastore. Can be an ESX host, cluster or datacenter.
goldVolume	String	Name of volume used when creating more than one NFS-based datastore. This volume is not permanently presented to the ESX hosts. This volume becomes the parent of the FlexClones.

Type	Value	Description
protocol	String	The protocol being used. Valid values are NFS, FCP, iSCSI.
containerName	String	Name of the aggregate for new NFS datastores, or name of volume for new VMFS datastores.
sizeInMB	Long	Size of the datastore in MB. An additional 256 MB is added for VMFS datastores to cover metadata overhead.
thinProvision	Boolean	If true, space will not be reserved for the storage object. For NFS, the volume will guarantee=none . For VMFS, LUN will be created with '-o noreserve'.
volAutoGrow	Boolean	If true, the volAutoGrowInc and volAutoGrowMax values are applied to the volume.
volAutoGrowInc	Long	Increment in which to grow volume automatically as needed in MB.
volAutoGrowMax	Long	Maximum size to which to grow the volume automatically in MB.
datastoreNames	List <String>	List of datastore names. Care should be taken by the application to prevent duplicate datastore, volume or LUN names. For NFS, datastore name is used as volume name. For VMFS, datastore name is used as LUN name.
numDatastores	Int	Number of datastores. This should indicate the size of the list of names in datastoreNames .
blockSize	Int	VMFS block size in MB.
controller	ControllerSpec	The controller.
wrapperVol	Boolean	When true, new volume is created to contain the new LUN to be used for a new VMFS datastore (containerName must contain aggregate name if true). When true, the volume containing the LUN (VMFS datastore) will be resized to make room for the new size of the LUN (if required).

Required parameters:

Some actions require the use of multiple parameters.

Specifying an existing datastore

- **mor**
- **controller**

Specifying new NFS datastores using createClones or createDatastore

- **targetMor** - only required for createDatastore
- **containerName**
- **sizeInMB**
- **thinProvision**
- **volAutoGrow**
- **volAutoGrowInc**
- **volAutoGrowMax**
- **protocol** - must be NFS
- **controller**
- **datastoreNames** - only one name in the list
- **numDatastores** - should be 1

Specifying new NFS datastores using createClones

- **goldVolume**
- **containerName**
- **sizeInMB**
- **thinProvision**
- **volAutoGrow**
- **volAutoGrowInc**
- **volAutoGrowMax**
- **protocol** - Must be NFS
- **controller**
- **datastoreNames**
- **numDatastores**

Specifying new VMFS datastores using createClones or createDatastore

- **targetMor** - only required for createDatastore
- **containerName**
- **sizeInMB**
- **thinProvision**
- **protocol** - must be FCP or iSCSI
- **controller**
- **datastoreNames** - only one name in the list
- **numDatastores** - should be 1

Specifying new VMFS datastores using createClones

- **containerName**
- **sizeInMB**

- **thinProvision**
- **protocol** - must be FCP or iSCSI
- **controller**
- **datastoreNames**
- **numDatastores**

XML

```

<complexType name="datastoreSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="containerName" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="blockSize" type="{http://www.w3.org/2001/XMLSchema}int"
          minOccurs="0"/>
        <element name="controller" type="{http://server.kamino.ibm.com/}controllerSpec"
          minOccurs="0"/>
        <element name="datastoreNames" type="{http://www.w3.org/2001/XMLSchema}string"
          maxOccurs="unbounded" minOccurs="0"/>
        <element name="goldVolume" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="mor" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="numDatastores" type="{http://www.w3.org/2001/XMLSchema}int"/>
        <element name="protocol" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="sizeInMB" type="{http://www.w3.org/2001/XMLSchema}long"
          minOccurs="0"/>
        <element name="targetMor" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="temp" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="thinProvision" type="{http://www.w3.org/2001/XMLSchema}boolean"/>
        <element name="volAutoGrow" type="{http://www.w3.org/2001/XMLSchema}boolean"/>
        <element name="volAutoGrowInc" type="{http://www.w3.org/2001/XMLSchema}long"
          minOccurs="0"/>
        <element name="volAutoGrowMax" type="{http://www.w3.org/2001/XMLSchema}long"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

ControllerSpec

The ControllerSpec specification describes the controller connection and authentication data. This information is used by the Provisioning and Cloning capability to connect to the controller using the ZAPI interface. No other protocol is used to connect to the controller.

Properties

Type	Value	Description
ipAddress	String	IP or host name of the controller.
username	String	User name (does not need to be root).
password	String	Password.

Type	Value	Description
ssl	Boolean	If true, use HTTPS. If false, use HTTP to connect to the controller.
passthroughContext	String (optional)	Name of the vFile or Vserver on which to create the new storage.
actuallyOpsMgr	Boolean (optional)	If true, connects to Operations Manager. If false, connects to controller.

XML

```
<complexType name="controllerSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="ipAddress" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="password" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="ssl" type="{http://www.w3.org/2001/XMLSchema}boolean"/>
        <element name="username" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

VmSpec

The VmSpec specification describes configuration and action options for each virtual machine created.

The MAC address must be in a range defined by VMware. Refer to VMware documentation for more information.

Properties

Type	Value	Description
macAddress	Map <String, String>	Virtual network adapter to MAC address information (optional).
custSpecName	String	Name of guest customization specification to be applied (optional).
vmMoref	String	The managed object reference of the virtual machine to be redeployed.
powerOn	Boolean	If true, the new virtual machines are powered on after all have been created.

XML

```
<complexType name="vmSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
```



```

<element name="custSpec" type="{http://server.kamino.ibm.com/}guestCustomization
Spec"
minOccurs="0"/>
<element name="domain" type="{http://server.kamino.ibm.com/}domainSpec"
minOccurs="0"/>
<element name="macAddress">
  <complexType>
    <complexContent>
      <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
        <sequence>
          <element name="entry" maxOccurs="unbounded" minOccurs="0">
            <complexType>
              <complexContent>
                <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
                  <sequence>
                    <element name="key" type="{http://www.w3.org/2001/XMLSchema}string"
minOccurs="0"/>
                    <element name="value" type="{http://www.w3.org/2001/XMLSchema}string"
minOccurs="0"/>
                  </sequence>
                </restriction>
              </complexContent>
            </complexType>
          </element>
        </sequence>
      </restriction>
    </complexContent>
  </complexType>
</element>
<element name="powerOn" type="{http://www.w3.org/2001/XMLSchema}boolean"/>
<element name="vmMoref" type="{http://www.w3.org/2001/XMLSchema}string"
minOccurs="0"/>
</sequence>
</restriction>
</complexContent>
</complexType>

```

GuestCustomizationSpec

The GuestCustomizationSpec specification identifies the guest customization specification.

Properties

Type	Value	Description
name	String	Name of the guest customization specification.
useVmName	Boolean	If guest customization specification is of type CustomizationSysprepText, this option can be used to make the guest hostname match the virtual machine name. Note: It is the responsibility of the implementer to ensure that the virtual machine name results in a valid host name.

XML

```
<complexType name="guestCustomizationSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="name" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="useVmName" type="{http://www.w3.org/2001/XMLSchema}boolean"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

ConnectionBrokerSpec

The ConnectionBrokerSpec specification describes the connection broker and authentication data. This information is used by VSC for VMware vSphere to connect to the connection broker (VMware View or Citrix XenDesktop).

Properties

Type	Value	Description
type	ConnectionBrokerType	The type of connection broker to import into. This can be VMWARE_VIEW_4_0, VMWARE_VIEW_4_5, or XEN_DESKTOP.
host	String	Hostname or IP address of the connection broker (used only for View import).
username	String	A user who can access the View Server (used only for View import).
password	String	Password of the specified user (used only for View import).
domain	String	FQDN where the connection broker resides.
desktopType	DesktopType	INDIVIDUAL_DESKTOP or DESKTOP_POOL (used only for View import).
accessMode	AccessMode	PERSISTENT or NON_PERSISTENT. This corresponds to dedicated and floating in View 4.5 and higher, respectively (used only for View import).
poolType	PoolType	NEW or EXISTING. Create a new pool or use an existing one (used only for View import).
poolName	String	The name of the new pool if the PoolType is set to NEW (used only for View import).

StatusMessage

You can use the StatusMessage specification to obtain the progress and status of an operation.

Properties

Type	Value	Description
id	Int	Operation identifier.
progress	Int	Valid values are 0-100. Indicates how much of the copy or clone process has completed at the time of the query.
status	String	Valid values are complete (finished without error), failed (finished with error), or running (operation in progress).
reason	String	If the status is failed, this contains the reason for the failure.

XML

```
<complexType name="statusMessage">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="id" type="{http://www.w3.org/2001/XMLSchema}int"/>
        <element name="progress" type="{http://www.w3.org/2001/XMLSchema}int"/>
        <element name="reason" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
        <element name="status" type="{http://www.w3.org/2001/XMLSchema}string"
          minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

Provisioning and Cloning sample code

```
<ns2:createClones xmlns:ns2="http://server.kamino.ibm.com/">
  <arg0>
    <cloneSpec>
      <clones>
        <entry>
          <key>apiTestClone1</key>
          <value>
            <powerOn>>false</powerOn>
          </value>
        </entry>
        <entry>
          <key>apiTestClone2</key>
          <value>
            <powerOn>>false</powerOn>
          </value>
        </entry>
        <entry>
          <key>apiTestClone3</key>
          <value>
            <powerOn>>false</powerOn>
          </value>
        </entry>
      </clones>
    </cloneSpec>
  </arg0>
</ns2:createClones>
```

```

    </value>
  </entry>
</clones>
<containerMoref>Datacenter:datacenter-2</containerMoref>
<files>
  <destDatastoreSpec>
    <controller>
      <ipAddress>10.10.10.2</ipAddress>
      <password></password>
      <ssl>>false</ssl>
      <username>root</username>
    </controller>
    <mor>Datastore:datastore-17</mor>
    <numDatastores>0</numDatastores>
    <thinProvision>>false</thinProvision>
    <volAutoGrow>>false</volAutoGrow>
  </destDatastoreSpec>
  <sourcePath>[unitTestSourceNFS]demoSource/demoSource.vmx</sourcePath>
</files>
<files>
  <destDatastoreSpec>
    <controller>
      <ipAddress>10.10.10.2</ipAddress>
      <password></password>
      <ssl>>false</ssl>
      <username>root</username>
    </controller>
    <mor>Datastore:datastore-17</mor>
    <numDatastores>0</numDatastores>
    <thinProvision>>false</thinProvision>
    <volAutoGrow>>false</volAutoGrow>
  </destDatastoreSpec>
  <sourcePath>[unitTestSourceNFS]demoSource/demoSource.vmdk</sourcePath>
</files>
<templateMoref>VirtualMachine:vm-255</templateMoref>
</cloneSpec>
<serviceUrl>https://10.10.10.2/sdk</serviceUrl>
<vcPassword>pass123</vcPassword>
<vcUser>Administrator</vcUser>
</arg0>
</ns2:createClones>

```

Provisioning and Cloning client-side programming

Various client-side programming environments allow you to access the SOAP service. You can use this service for your own client-side programming.

Accessing the SOAP API through Java

You can access the SOAP API by using Java-based tools or the Java programming language.

Generating a certificate for use with the wsimport tool:

You must generate a certificate for use with wsimport so that the tool will be able to read the WSDL that is generated by the build process.

About this task

Complete the following steps to generate a certificate for use with wsimport.

Procedure

1. Stop the NVPF service. This step is optional if you have already generated an SSL certificate for the environment.

2. Run the following command in the VSC for VMware vSphere installation directory: **c:\Program Files\IBM\Virtual Storage Console>bin\vsc ssl setup -domain <domain>**. For *<domain>*, enter the host name of the system running VSC for VMware vSphere or a fully qualified domain name of the system running VSC for VMware vSphere.

Note: This step is optional if you have already generated an SSL certificate for the environment.

3. From the VSC for VMware vSphere installation directory, change to `\etc` and run the command: **keytool -export -alias nvpf -keystore nvpf.keystore -file nvpf.cer**

Note: If you moved the keystore file from the `c:\Program Files\IBM\Virtual Storage Console\etc` directory, enter the path to the keystore file.

This command creates a new file called `nvpf.cer`. This certificate will be imported to the local Java keystore. If you have the Java JRE version 1.6.0_21 installed in Program Files, the command to execute will look like this:

```
c:\Program Files\IBM\Virtual Storage Console\etc>keytool -import -alias nvpf -file nvpf.cer -keystore "c:\Program Files\Java\jdk1.6.0_21\jre\lib\security\cacerts"
```

4. Enter keystore password: `changeit`
5. Run **wsimport** to grab the WSDL and generate the Java classes to write your own client. Enter the commands: `cd \dev c:\dev\wsimport -verbose -s . -p com.ibm.kamino.api https://<domain>:8143/kamino/public/api?wsdl` Be sure to change *<domain>* to a valid host or domain name. If that name does not resolve through DNS, add it to your hosts file (not localhost).

Accessing SOAP through C#

To begin using the SOAP API, you must first add the web reference to your project.

After the web service reference has been added, you can start accessing the client side objects to make API calls into the SOAP service.

Troubleshooting

This section describes how to troubleshoot general VSC for VMware vSphere installation and usage issues.

Issues that apply to multiple capabilities

There are several troubleshooting suggestions and issues that apply to all of the VSC for VMware vSphere capabilities. This section lists those issues.

Check the Release Notes

The *Release Notes* contain the most up-to-date information on known problems and limitations. The *Release Notes* also contain information on how to look up information about known issues.

The *Release Notes* are updated when there is new information about VSC for VMware vSphere. It is a good practice to check the *Release Notes* before you install VSC for VMware vSphere and any time you encounter a problem with VSC for VMware vSphere.

VMware only supports selecting one object when using right-click actions

VMware limits how many objects can be selected using right-click actions when running a vCenter plug-in such as VSC for VMware vSphere. When you select multiple VMware objects using a right-click action, VMware performs the operation on the first object only.

For example, the Provisioning and Cloning capability uses a wizard-driven workflow to select a VMware object in the Inventory panel. If you use a right-click action to select multiple objects, VMware ignores all but the first object selected. This is why the Provisioning and Cloning capability only supports space reclamation on a single virtual machine (VM) or datastore.

This limitation applies only to VMware objects that are selected through the Inventory panel. If you are working with an object that is owned by VSC for VMware vSphere, such as an object owned by the Monitoring and Host Configuration capability, you can use a multi-select action.

Issues that apply to the Monitoring and Host Configuration capability

This section contains information on troubleshooting tips and issues that affect the Monitoring and Host Configuration capability.

Getting information about storage controllers with an Alert status

Display the Status Reason column or open FilerView to get an explanation of why a storage controller displays the Alert status on the Monitoring and Host Configuration capability Overview panel. By default, the Status Reason column is hidden.

Procedure

1. On the Monitoring and Host Configuration capability Overview panel, click the arrow in the column header and select **Columns**.
2. Select **Status Reason**. Note that vFiler units do not return a detailed status reason; This controller is a MultiStore vFiler unit is displayed for vFiler units. The displayed text is truncated if the status reason is too long, and ellipses (...) are appended to the displayed text.
3. To get more information, right-click the storage controller and select **Open FilerView**. The FilerView GUI for the controller opens in your web browser.

Getting information about an ESX and ESXi host with an Alert status

You can display the causes of Alert icons for ESX hosts on the Monitoring and Host Configuration capability Overview panel. There are several options for learning more about an Alert.

About this task

The Status Reason column contains an explanation of why an Alert is being displayed. By default, the Status Reason column is hidden.

Also, the Monitoring and Host Configuration capability displays an Alert icon in the **Status** column for a host if vCenter reports alarms for that host.

In addition, the Monitoring and Host Configuration capability displays an Alert icon the **Adapter Settings**, **MPIO Settings**, or **NFS Settings** columns if a current host setting is different than the supported value set by the capability.

Procedure

To check the Alert status, take one of the following actions:

Option	Description
To check the status of the host	Display the Status Reason column. <ol style="list-style-type: none">1. On the Monitoring and Host Configuration capability Overview panel, click the arrow in the column header and select Columns.2. Select Status Reason.
For an Alert in the Status column	See the Alarms tab in vCenter for the host.
For an Alert in the other columns on the Monitoring and Host Configuration capability Overview panel	Right-click the ESX or ESXi host and select Show Details . Compare the host settings with the values set by the Monitoring and Host Configuration capability. For NFS settings, the current value is shown in red if it does not match the recommended value, and the recommended value is displayed. For some settings, the default is used and no explicit value is set. These default values are shown as empty strings ("").

What to do next

Correct any alarms reported for the host. Select **Set Recommended Values** from the right-click menu to have the Monitoring and Host Configuration capability update the desired ESX or ESXi host settings.

Related reference:

"ESX host settings set by Monitoring and Host Configuration capability" on page 27

Collecting the VSC for VMware vSphere log files

You can collect the log files from all installed VSC for VMware vSphere components and capabilities using the Data Collection panel in the Monitoring and Host Configuration capability. Technical support might ask you to collect the log files to help troubleshoot a problem.

Procedure

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the Inventory panel, and then select the IBM N series tab.
3. In the Monitoring and Host Configuration capability, select the Data Collection panel.
4. Select **Export VSC Logs** and click **Submit**.
5. When prompted, save the file to your local computer.

What to do next

Send the .zip file to technical support.

Troubleshooting error message "The client cannot communicate with the Virtual Storage Console Server"

The vSphere Client displays the message "The client cannot communicate with the Virtual Storage Console Server. Verify that the server is running and that the client has network connectivity to the server." Correct any connectivity problems, and if necessary restart the affected servers.

Procedure

1. If the VSC for VMware vSphere and vCenter servers run on separate systems, verify that the two systems have basic network connectivity between them.
2. Verify that any firewalls on the server allow communication on TCP port 443 (HTTPS/SSL).
3. Verify that the vCenter account credentials used for background discovery have not expired.
4. Log into the Windows server or servers verify that the VSC for VMware vSphere service is running and the VMware vCenter services are running.
5. If a service is not running, restart the service.
6. If none of the previous steps solved the problem, reboot the Windows server or servers running VSC for VMware vSphere and vCenter Server.

Related tasks:

“Updating vCenter credentials for background discovery”

Updating vCenter credentials for background discovery

If the vCenter credentials specified when VSC for VMware vSphere was installed expire, the Monitoring and Host Configuration capability is no longer able to run background discovery tasks. The Monitoring and Host Configuration capability displays an error message. Re-register VSC for VMware vSphere to enter updated credentials.

Before you begin

The vCenter account must be an administrator-level account.

Procedure

1. Click the link in the error message about expired credentials, or point a Web browser to the registration Web page: `https://hostname:8143/Register.html`
hostname is the host name or IP address of the server where VSC for VMware vSphere is installed.
If a security certificate warning is displayed, choose the option to ignore it or to continue to the Web site. The plug-in registration Web page is displayed with the current credentials.
2. Enter the new password for the user name shown, or enter a new user name and password.
3. Restart all vCenter Clients.

Resolution of issues with the Backup and Recovery capability

If you encounter unexpected behavior while configuring the Backup and Recovery capability or during a backup or restore operation, you can follow specific troubleshooting procedures to identify and resolve the cause of such issues.

Backup and Recovery capability values that you can override

To improve operational efficiency, you can modify the `smvi.override` configuration file to override the default values specified in the file. These values control such settings as the number of VMware snapshots that are created or deleted during a backup or the amount of time before a backup script stops running.

The `smvi.override` configuration file is located under the installation directory at `C:\Program Files\IBM\VSC\smvi\server\etc\smvi.override`.

If you modify any of these entries, you must restart the server for the changes to take effect. You can modify the following entries if you need to override the default values:

vmware.max.concurrent.snapshots=6

This entry specifies six as the default maximum number of VMware snapshots created or deleted per datastore during a backup.

vmware.quiesce.retry.count=0

This entry specifies zero as the maximum number of retry attempts for VMware snapshots.

vmware.quiesce.retry.interval=5

This entry specifies the amount of time, in seconds, between retry attempts for VMware snapshots.

vim.client.log.verbose=true

This entry, when the value is true, logs the interactions between the SMVI server and the vCenter server.

smvi.script.timeout.seconds=120

This entry specifies the SMVI timeout value for a pre-backup or post-backup script, which is when the SMVI server stops waiting for the script to finish running.

Backup and Recovery capability event and error logs

The Backup and Recovery capability logs both server messages and messages between the server and the user interface, including detailed information about event messages and errors. Reviewing these logs helps you troubleshoot any errors that occur during backup or restore operations.

The log files are stored under the installation directory at the following locations:

- The server log messages are at C:\Program Files\IBM\Virtual Storage Console\smvi\server\log\server.log.
- The log messages between the user interface and the server are at C:\Program Files\IBM\Virtual Storage Console\log\smvi.log.

Email notification for scheduled backup contains a broken link

Issue

When you click the link to view the log files in the email notification for a backup job, an error occurs.

Cause

This problem occurs if you disable the IP address of the network adapters for the SMVI log viewer.

Corrective action

You must enable the IP address of the network adapters for the SMVI log viewer in one of the following ways, depending on your operating system:

- Select **Control Panel > Network connections > Network and Sharing Center** in Windows 2003, Windows 2008, Windows 2008 R2, and Windows 7 environments.
- Select **Control Panel > Network connections** in Windows Vista and Windows XP environments.

You may have reached the maximum number of NFS volumes configured in the vCenter

Message

You may have reached the maximum number of NFS volumes configured in the vCenter. Check the vSphere Client for any error messages.

Description

This message occurs when you attempt to mount a backup of an NFS datastore on a Vserver with the root volume in a load-sharing mirror relationship and the mount fails.

Corrective action

When you add a storage system operating in Cluster-Mode to the Backup and Recovery capability, use the server's IP address instead of the Vserver IP address.

Using ESX hosts with IBM N series storage

To use ESX hosts with IBM N series storage systems, you need to correctly provision storage and configure ESX hosts.

The following sections provide general storage system and ESX host information.

LUN type guidelines

LUNs must be created with the correct LUN type.

If the LUN will be configured with VMFS, then use the LUN type `vmware`.

If the LUN will be configured with RDM, then use the guest OS for the LUN type. For more information on the LUN type to use, see the *Data ONTAP Block Access Management Guide for iSCSI and FC* for your version of Data ONTAP software.

Manually provisioning storage

To configure your storage systems to connect to virtual machines (VMs) running on VMware ESX or ESXi, you must create new volumes, LUNs, and igroups and then map the LUNs to the igroups.

Before you begin

You need the FC or iSCSI identifiers of the ESX or ESXi host.

For detailed instructions on how to perform the following steps, see the *SAN Administration Guide* (called *Block Access Management Guide for iSCSI and FC* in Data ONTAP 8.1 and earlier) for your version of Data ONTAP software.

About this task

The Provisioning and Cloning capability can also be used to provision storage.

Procedure

1. Create an initiator group (igroup) for each VMware ESX or ESXi server using the `vmware` igroup type. Use the WWPNs for all FC HBAs in the ESX or ESXi host or the iSCSI initiator node name of the ESX or ESXi host when creating the igroup.

For ESX 4.0, 4.1, and 5.0 configurations that support ALUA, enable the ALUA option on the igroup.

Note: (Data ONTAP operating in 7-Mode) If your storage systems are running Data ONTAP operating in 7-mode and you are using a Microsoft cluster (Windows Server 2003 MSCS or Server 2008 failover cluster) configuration, do not enable ALUA for LUNs used by guest operating systems. If ALUA is enabled, the cluster loses its persistent reservations during storage faults, causing the cluster service to be unavailable. In addition, you need to avoid ALUA mismatches so that no initiator participates in both ALUA and non-ALUA enabled groups.

Check the N series Interoperability Matrices website (accessed and navigated as described in Websites) at www.ibm.com/systems/storage/network/interophome.html to see which versions of Data ONTAP are supported with Microsoft cluster (Windows Server 2003 MSCS and Server 2008 failover cluster) configurations.

2. Create the storage for each virtual machine.
 - a. Create one or more volumes to contain the LUNs. FlexVol volumes are recommended in general, and are required if you are using Snapshot copies.
 - b. Create a LUN for the VM's root disk.
 - c. Create any additional LUNs needed for the VM's application data.
 - d. Map all of the LUNs to the igroup for the ESX or ESXi host. If you plan to use VMotion to move your guest operating systems from one VMware ESX or ESXi host to another, map the LUN to all hosts in the cluster. The LUN IDs must be identical.
3. Optionally, verify, and if necessary correct the alignment of the VMDK partitions.

Sometimes partition misalignment problems can arise, which can lead to performance degradation under very heavy I/O. Depending on your configuration, you might need to align your VMDK partitions to avoid subsequent performance problems.

If you use RDM, and you use the correct guest OS for the LUN type, you should not experience alignment problems.


How to set up VMware ESX

After creating the necessary LUNs and igroups, and after mapping the LUNs to the igroups, you must configure your host.

For guests in a Microsoft Windows cluster (MSCS) configuration only, you also need to change the path selection policy.

For more information about setting up ESX or ESXi, see the *SAN Configuration Guide* (called *Fibre Channel and iSCSI Configuration Guide* in Data ONTAP 8.1 and earlier) and the VMware product documentation for your version of ESX.

Related information:

 VMware product documentation - www.vmware.com/support/pubs/vi_pubs.html

Configuring the VMware ESX host

Configuring the VMware ESX or ESXi host requires rescanning the bus, creating a datastore, and creating a new VM.

Before you begin

You need to create the required LUNs before starting this task.

More information on this task is available in the documentation provided by VMware.

Procedure

1. Rescan the SCSI bus to discover the new LUNs.
 - a. Open the VMware vSphere Client and connect to your ESX or ESXi host.

- b. On the **Configuration** tab, select **Hardware > Storage Adapters**
- c. Click **Rescan**.
2. Create a VMFS datastore on the LUN.
3. Create a new VM or add a new disk to an existing VM.

(Data ONTAP operating in 7-Mode) Manually setting the path selection policy for Microsoft cluster configurations

If the storage system is running Data ONTAP operating in 7-mode and you are using a Microsoft cluster (Windows Server 2003 MSCS or Server 2008 failover cluster) configuration, you should disable ALUA on the igroup and change the path selection policy to FIXED for guest operating systems.

About this task

For Microsoft Windows guest operating systems in a cluster configuration, always use the FIXED path selection policy and disable ALUA on the igroup for the LUNs. This might require you to manually set the path selection policy.

If ALUA is enabled, the Windows cluster loses its persistent reservations during storage faults, causing the cluster service to be unavailable. When ALUA is disabled, the FIXED path selection policy is required to avoid sending I/O over proxy paths.

The default path selection policy set by VSC for VMware vSphere should be used if the ESX or ESXi host does **NOT** have guest operating systems in a Windows cluster (MSCS or failover cluster) configuration. For ESX 4.0, 4.1, and 5.0 systems, the path selection policy is set to round robin (RR) for ALUA FC configurations and all iSCSI configurations, and set to FIXED for non-ALUA configurations. For ESX 3.5, the default policy is FIXED.

Procedure


1. To manually change the path selection policy, enter the following command on the ESX or ESXi host:


```
esxcli conn_options nmp device setpolicy --device device_name --psp VMW_PSP_FIXED
```

 For more information about the **esxcli** command, see the *vSphere Command-Line Interface Installation and Reference Guide*.
2. To disable ALUA for an igroup, enter the following command at a Data ONTAP command prompt:


```
igroup set igroup_name alua no
```

Related information:

 [vSphere Command-Line Interface Installation and Reference Guide - www.vmware.com/pdf/vsphere4/r40/vsp_40_vcli.pdf](http://www.vmware.com/pdf/vsphere4/r40/vsp_40_vcli.pdf)

Timeout values for guest operating systems

The guest OS (GOS) timeout scripts set the SCSI I/O timeout values for supported Linux, Solaris, and Windows guest operating systems. The timeout values ensure correct failover behavior.

These scripts are provided as .ISO files in the Tools panel of the Monitoring and Host Configuration capability. There are two scripts for each operating system:

- A 60-second script

- A 190-second script

In most cases, the recommended value is 60 seconds.

You can mount and run the script from the vSphere client. The Tools panel provides URLs for the scripts.

To get the script containing the timeout values you want for your operating system, you must copy the correct URL from the Tools panel and mount it as a virtual CD-ROM in the virtual machine using the vSphere client. Make sure you install the script from a copy of VSC for VMware vSphere that is registered to the vCenter Server that manages the virtual machine. After the script has been installed, you can run it from the console of the virtual machine.

The section *Installing GOS scripts* contains for detailed steps for performing this task.

Running the GOS timeout scripts for Linux

The guest operating system timeout scripts set the SCSI I/O timeout settings for RHEL4, RHEL5, RHEL6, SLES9, SLES10, and SLES11. You can specify either a 60-second timeout or a 190-second timeout. You should always run the script each time you upgrade to a new version of Linux.

Before you begin

You must mount the ISO image containing the Linux script before you can run it in the virtual machine.

Procedure

1. Open the console of the Linux virtual machine and log in to an account with root privileges.
2. Run the `linux_gos_timeout-install.sh` script.

Results

For RHEL4 or SLES9, a message similar to the following is displayed:

```
Restarting udev... this may take a few seconds.  
Setting I/O Timeout (60s) for /dev/sda - SUCCESS!
```

For RHEL5 or RHEL6, a message similar to the following is displayed:

```
patching file /etc/udev/rules.d/50-udev.rules  
Hunk #1 succeeded at 333 (offset 13 lines).  
Restarting udev... this may take a few seconds.  
Starting udev: [ OK ]  
Setting I/O Timeout (60s) for /dev/sda - SUCCESS!
```

For SLES10 or SLES11, a message similar to the following is displayed:

```
patching file /etc/udev/rules.d/50-udev-default.rules  
Hunk #1 succeeded at 114 (offset 1 line).  
Restarting udev ...this may take a few seconds.  
Updating all available device nodes in /dev:           done
```


What to do next

Unmount the ISO image by clicking the **CD/DVD Connections** icon in the vSphere Client and selecting **CD/DVD Drive 1 > Disconnect from filename.iso**.

Running the GOS timeout scripts for Solaris

The timeout scripts set the SCSI I/O timeout settings for Solaris 10. You can specify either a 60-second timeout or a 190-second timeout.

Before you begin

You must mount the ISO image containing the Solaris script before you can run it in the virtual machine.

Procedure

1. Open the console of the Solaris virtual machine and log in to an account with root privileges.
2. Run the `solaris_gos_timeout-install.sh` script.

Results

For Solaris 10, a message similar to the following is displayed:
Setting I/O Timeout for /dev/s-a - SUCCESS!

What to do next

Unmount the ISO image by clicking the **CD/DVD Connections** icon in the vSphere Client and selecting **CD/DVD Drive 1 > Disconnect from filename.iso**.

Running the GOS timeout script for Windows

The timeout scripts set the SCSI I/O timeout settings for Windows guest operating systems. You can specify either a 60-second timeout or a 190-second timeout. You must reboot the Windows guest OS for the settings to take effect.

Before you begin

You must mount the ISO image containing the Windows script before you can run it in the virtual machine.

Procedure

1. Open the console of the Windows virtual machine and log in to an account with Administrator privileges.
2. If the script does not automatically start, open the CD drive and run `windows_gos_timeout.reg`. The Registry Editor dialog is displayed.
3. Click **Yes** to continue. The following message is displayed:
The keys and values contained in D:\windows_gos_timeout.reg have been successfully added to the registry.
4. Reboot the Windows guest OS.

What to do next

Unmount the ISO image by clicking the **CD/DVD Connections** icon in the vSphere Client and selecting **CD/DVD Drive 1 > Disconnect from filename.iso**.

How to identify and fix VMDK partition alignment issues

In some cases, VMDK partitions can become misaligned, leading to performance degradation.

For more information about partition alignment, see *Technical report: Storage Block Alignment with VMware Virtual Infrastructure and IBM System Storage N series - Correctly aligning LUNs and virtual disk files*.

Note: This issue may be serious enough to warrant action, but the performance degradation depends on your I/O load and configuration. In many cases, the decrease in performance will be negligible.

Also note that this problem is not unique to IBM N series storage platforms.

VMDK partitions need to be aligned at both the VMFS and guest OS levels. For example, you can align the partitions at the VMFS level by selecting the vmware LUN type when creating your LUNs. By doing so, the partitions are aligned to sector 128 or sector 0, depending on whether you use vCenter or vmkfstools to create the VMFS. Regardless, the partitions will be aligned as both are multiples of 4 KB, thereby fulfilling the WAFL read/write requirements.

However, because some operating systems implement sector offsets that are not aligned to 4 KB boundaries, the partitions might still not be aligned at the guest OS level. Therefore, you must manually align the .vmdk files at the guest OS level for VMFS and NFS datastores.

Note: If you use RDM and create the LUN with the correct guest OS for the LUN type, then you should not experience alignment issues with the RDM LUNs. The base VMDK might still have an alignment issue.

Related information:



Storage Block Alignment with VMware Virtual Infrastructure and IBM System Storage N series - <ftp://service.boulder.ibm.com/storage/isv/NS3593-0.pdf>

Checking VMDK partition alignment with mbralign

You can use either the Optimization and Migration capability or the mbralign tool included with VSC for VMware vSphere to check VMDK partition alignment.

About this task

The mbralign tool replaces the mbrscan tool and is effective on -flat.vmdk and fixed .vhd files that are partitioned using a master boot record (MBR) partition table.

Note: There is a version mbralign for ESX hosts and one for ESXi hosts. When you download mbralign, you choose the software package based on whether you have an ESX host or an ESXi host. The example in these steps uses the standard ESX mbralign. The mbralign for ESXi does not support all the features used here, such as drive letter restoration.

If you do not want to power down the VM, take one of the following actions:

- Use the Optimization and Migration capability online alignment tool.

Note: For information on using the Optimization and Migration capability tool to scan virtual machines (VM) and perform online alignments, see the section on Optimization and Migration.

- Use `mbralign` and take either a Data ONTAP Snapshot copy of the volume containing the datastore LUN or NFS datastore or a VMware snapshot of the VM in question.

Then run `mbrscan` against the copy.

Procedure

1. On the ESX host console, change to the directory where `mbralign` is installed.
2. Enter the following command on the ESX host console: `mbralign { --scan all | filename }`

The `--scan all` option scans all `-flat.vmdk` files.

filename specifies the name of a single file to scan. The command displays information indicating whether the VMDK partition is correctly aligned.

Example

```
# /opt/ontap/santools/mbralign --scan all
Building file list...
/vmfs/volumes/4c604abb-e41943c0-a81f-001b7845166c/win2k3sp2_64v_esx-09/
win2k3sp2_64v_esx-09_1-flat.vmdk P1 lba:63 Aligned: No

/vmfs/volumes/4c604abb-e41943c0-a81f-001b7845166c/win2k3sp2_64v_esx-09/
win2k3sp2_64v_esx-09_2-flat.vmdk P1 lba:63 Aligned: No
```

VMDK partition alignment with mbralign overview

The `mbralign` tool enables you to correct misaligned VMDK partitions.

Note: The Optimization and Migration capability allows you to perform online alignments. In addition, you can also use the VMware vCenter Converter to perform offline alignments.

The `mbralign` tool works on primary VMDK partitions with a master boot record (MBR) partition table. If there are multiple partitions, the partitions must be in order on the disk.

Starting with the `mbralign` tool in VSC 2.1 for VMware vSphere, the 1-TB size limit on VMDKs has been removed. You can use `mbralign` with any VMDK size supported by VMware.

The `mbralign` tool has the following requirements:

- The destination datastore must have enough free space for a full copy of the `-flat.vmdk` file.
- GRUB-booted Linux guest operating systems need to have GRUB reinstalled after aligning the boot partition.
- The virtual machine using the VMDK must be shut down when running `mbralign`. If you use the feature to preserve Windows drive letter mapping, the `mbralign` program shuts down the VM after collecting drive letter information.
- For ESX clusters, you must run the `mbralign` program on the ESX or ESXi host where the VM is currently registered. For NFS datastores, the `mbralign` program cannot detect if the VM is powered down if the VM is running on another ESX or ESXi host.

The `mbralign` tool has the following limitations:

- Large NFS filesystems mounted on the ESX or ESXi host can greatly increase the time required to run `mbralign`, because `mbralign` scans them for VMDKs. Temporally unmounting large filesystems that do not contain VMDKs needing alignment should improve performance.
- VMDKs containing Windows dynamic disks or GPT partitions are not supported. The Windows operating system must be installed on the C: drive.
- VMDKs containing Linux LVM are not supported.
- The `mbralign` tool does not fix misalignment associated with linked clone or VMware delta files; it can only fix the base VMDK. All VMware snapshots and linked clones must be removed from the VM using the disk being aligned.
- If the storage system volume containing the VMDK is included in Snapshot copies, the alignment process can cause the Snapshot copies to grow very large (up to twice the size of the volume). The space used for the Snapshot copies is not released when the backup file is deleted. Deleting the Snapshot copies from before the alignment process releases the space used.
- Media devices, such as CD-ROM or DVD drives used by the VM, might not map to their Windows original drive letters after running the `mbralign` program. This can happen when there are multiple media drives or when the drive contains media.
- Do not use the `--force` option of the `mbralign` command on VMDKs for virtual machines running Windows 7, Windows Server 2008, or Windows Server 2008 R2. This can corrupt the boot LUN.
- Do not use `mbralign` with Solaris guest operating systems; it cannot correctly align them. Solaris ZFS file systems should not have alignment issues.

Starting with the `mbralign` tool in VSC 2.1 for VMware vSphere, you can now preserve the original drive mapping of Windows disks. Earlier versions of `mbralign` could only ensure that the C:\ drive mapped to the correct partition.

- The Windows virtual machine must be running when you start `mbralign` so that `mbralign` can collect the drive letter mapping information. The `mbralign` program prompts lead you through the process of shutting down the VM after collecting drive letter information and then starting the actual alignment process.
- The Windows operating system folder must be on the C:\ drive. For example, C:\Windows.
- The VMware tools package must be installed on the VM. Be sure to use the version of VMware tools that matches the ESX version on which the VM is running. See your VMware documentation for instructions on installing the VMware tools in the guest operating system.
- For Windows Server 2000, you must install the Windows 2000 Resource Kit Tools for administrative tasks, which includes the `diskpart` program. Be sure to take the default installation location.
- For 64-bit Windows Server 2003 and Windows XP guests, install Windows hotfix KB 942589 on the VM.
- Note that this drive letter mapping process does not apply to Linux virtual machines.

You can use the `--bs` option of the `mbralign` command to increase the default 8 KB block size to a larger value (16, 32, 64, 128, or 1024) using the `--bs` option of the `mbralign` command. Doing this improves performance of VMFS datastores. A larger block size generally means that there are fewer reads and writes.

To see all of the command options, you can use `mbralign --help` command, or you can refer to the `mbralign` man page.

Related information:

-  Windows 2000 Resource Kit Tools for administrative tasks.
-  Hotfix KB 942589

Fixing VMDK partition alignment using mbralign

If a VMDK partition is misaligned, you can align the partition using the **mbralign** tool included with VSC for VMware vSphere.

Before you begin

Be aware that **mbralign** can take anywhere from 1 or 2 minutes to several minutes per gigabyte of storage in the affected files.

Note: The Optimization and Migration capability allows you to perform online alignments. In addition, you can use the VMware vCenter Converter to perform offline alignments.

Procedure

1. Remove any VMware snapshots from the VM that is to be realigned.
2. Temporarily unmount large NFS filesystems that do not contain VMDKs needing alignment from the ESX or ESXi host.
3. Shut down the VM.
4. For Linux VMs, and Windows VMs with only a C:\ drive, shut down the VM. For a Windows VM with multiple drive letters mapped, the VM must be running so that **mbralign** can collect the drive letter information.
5. On the ESX or ESXi host console, change to the directory containing the .vmdk file for the VM.
6. Enter the following command: `path/mbralign name.vmdk`
path is the path where the **mbralign** program is installed.
name is the name of the VMDK file being aligned.
7. If prompted, enter `yes` for a Windows VM to automatically collect and restore drive letters. Enter the Windows Administrator credentials for the VM. The VM is automatically shut down after the drive letter information is collected.
8. When prompted Are you sure that no snapshots/linked clones exist for this vmdk? Enter `y`.
Attention: The use of **mbralign** on a VMDK file that has a snapshot or linked clone associated with it can result in unrecoverable data loss or data corruption.
9. For Windows guest operating systems for which you are not using the drive letter restore feature, restart the VM and verify that the guest operating system boots successfully.
10. For Linux guest operating systems using the GRUB boot loader, reinstall GRUB before restarting the VM.
11. After verifying the VM has booted and is operating correctly, delete the backup files created by **mbralign**. These files are saved in the same directory as the .vmdk file and have names ending in `-mbralign-backup`.

Example

In the following example, some output has been deleted for clarity and the lines have been truncated to fit the page.

```

[root@esxhost1 VM2]# /opt/ontap/santools/mbralign VM2_1.vmdk
The vmdk file looks like it belongs to a Windows Virtual Machine: VM2.
Would you like to automatically align the vmdk and restore the original drive letters?
If this is not a Windows Virtual Machine, or if this is a VM part of a Microsoft Cluster,
select no (yes/no) yes
This VM also has the following vmdk files associated to it that also need to be aligned.
/vmfs/volumes/4bb1f98a-a2c428cc-f253-001e4f2f3dd3/VM2/VM2.vmdk
/vmfs/volumes/4bb1f98a-a2c428cc-f253-001e4f2f3dd3/VM2/VM2_2.vmdk
Do you want to also align /vmfs/volumes/4bb1f98a-a2c428cc-f253-001e4f2f3dd3/VM2/VM2.vmdk?
(yes/no/all) yes
Do you want to also align /vmfs/volumes/4bb1f98a-a2c428cc-f253-001e4f2f3dd3/VM2/VM2_2.vmdk?
(yes/no/all) yes
Checking the power state of the VM.....ON
Please provide Administrator credentials or [enter] to skip:
ESX Server Username: root
ESX Server Password:
VM Domain:
VM Username: Administrator
VM Password:
...
Alignment complete for VM2_1.vmdk
The next vmdk to align is: /vmfs/volumes/4bb1f98a-a2c428cc-f253-001e4f2f3dd3/VM2/VM2_2.vmdk
...
Press enter when you are ready to power on the VM [enter]
Powering on the VM.....DONE
Establishing connection to VM.....DONE
Collecting volume information.....DONE
Setting drive letters.....DONE
Removing temporary files.....DONE

```

```

[root@esxhost1 V_M_2]# /opt/ontap/santools/mbralign V_M_2.vmdk
Part Type old LBA New Start LBA New End LBA Length in KB
P1 83 63 64 208846 104391
P2 8e 208845 208856 16771871 8281507

NOTICE:
This tool does not check for the existence of Virtual Machine snapshots or linked clones.
The use of this tool on a vmdk file that has a snapshot or linked clone associated with it
can result in unrecoverable data loss and/or data corruption.
Are you sure that no snapshots/linked clones exist for this vmdk? (y/n)y

Creating a backup of V_M_2.vmdk
Creating a backup of ./V_M_2-flat.vmdk
Creating a copy the Master Boot Record
Working on partition P1 (2): Starting to migrate blocks from 32256 to 32768.
...

```

Reinstalling GRUB for Linux guests after running mbralign

After running **mbralign** on disks for Linux guest operating systems using the GRUB boot loader, you must reinstall GRUB to ensure that the guest operating system boots correctly.

Before you begin

The **mbralign** program has completed on the on the .vmdk file for the virtual machine.

About this task

This topic applies only to Linux guest operating systems using the GRUB boot loader and SystemRescueCd.

Procedure

1. Mount the ISO image of Disk 1 of the installation CDs for the correct version of Linux for the virtual machine.
2. Check the box for **Connected** (if the VM is running) or **Connect at power on** (if the VM is not running).
3. Open the vSphere Client remote console for the VM.
4. If the VM is running and hung at the GRUB screen, click in the display area to make sure it is active, then press Ctrl-Alt-Insert to reboot the VM. If the VM is not running, start it, and then immediately click in the display area to make sure it is active.
5. As soon as you see the VMware BIOS splash screen, press the Escape key once. The boot menu is displayed.
6. At the boot menu, select CD-ROM.
7. At the Linux boot screen, enter :linux rescue
8. Take the defaults for Anaconda (the blue/red configuration screens). Networking is optional.
9. Launch GRUB by entering: grub
10. If there is only one virtual disk in this VM, or if there are multiple disks, but the first is the boot disk, then run the following GRUB commands:
root (hd0,0) setup (hd0) quit

If you have multiple virtual disks in the VM, and the boot disk is not the first disk, or you are fixing GRUB by booting from the misaligned backup VMDK, enter the following command to identify the boot disk: find /boot/grub/stage1 Run the following commands:
root (boot_disk,0) setup (boot_disk)

quit *boot_disk* is the disk identifier of the boot disk.

11. Press Ctrl-D to log out. Linux rescue shuts down and then reboots.

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:
www.ibm.com/storage/nas
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:
www.ibm.com/storage/support/nseries/
This web page also provides links to AutoSupport information as well as other important N series product resources.
- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:
www.ibm.com/systems/storage/network/interophome.html
- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:
<http://publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp>

Copyright and trademark information

This section includes copyright and trademark information, and important notices.

Copyright information

Copyright ©1994 - 2013 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2011, 2013 IBM Corporation. All rights reserved.

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

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by

NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered 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.

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp is a licensee of the CompactFlash and CF Logo trademarks.

NetApp NetCache is certified RealSystem compatible.

Notices

This information was developed for products and services offered in the U.S.A.

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 on 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 give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

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 states 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 web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

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

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may

vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Index

Special characters

- /etc/smvi.override file
 - changing the SnapManager for Virtual Infrastructure server IP address 110
- .bat 164
- .cmd 164
- .sfr file
 - manually creating 109
 - Restore Agent
 - creating a .sfr file for 109
 - Single File Restore file
 - See .sfr file

A

- Access Management Console 66
- accounts
 - configuring with RBAC 13
- active backup job, suspending 104
- actual alignments 88
- Add Single File Restore Session wizard
 - creating a limited self-service restore session 113
 - creating a self-service restore session 111
- Admin Assisted
 - restore sessions 110
- administrator-assisted file restoration
 - of virtual disk files 108
- administrator-assisted restore sessions,
 - about 109
- aggregate 58, 67
 - adding to resources 67
- alert status
 - getting more information 168
- alignments 93
 - actual 88
 - functional 88
 - identifying and fixing partition 178
 - offline 88
 - online 88
- ALUA
 - enabling 173
- API
 - Provisioning and Cloning
 - methods 136
- API documentation
 - accessing 135
- APIs
 - for VMware vCloud 135
- APIs for Provisioning and Cloning 135
- APIs for VMware vCloud 135
- architecture
 - VSC for VMware vSphere 4
- authentication credentials
 - configuring the vCenter Server 115

B

- Backup and Recovery
 - adding or removing 16
 - lock management 6
 - SnapMirror requirements 97
 - Snapshot and SnapRestore
 - requirements 97
- Backup and Recovery capability 64
 - authentication credentials 98
 - configuration of 98
 - defined 97
 - network connection to Restore Agent 108
 - restoring data from backups 104
- backup jobs
 - deleting 103
 - troubleshooting 170
- backup mount error 171
- Backup Now window 103
- backup retention 100
- backup scripts 170
- Backup tab 112, 115
- Backup wizard 100, 101
- backups
 - adding 100
 - editing a backup job 103
 - failure 171
 - finding 106
 - managing 100
 - mounting 107
 - of datastores or datacenters 101
 - of virtual machines 100
 - restoring data from 104, 105
 - restoring data from backups with
 - failed VMware consistency snapshots 105
 - resuming a suspended backup job 104
 - selecting the restore destination 105
 - starting a one-time backup 103
 - suspending an active backup job 104
 - unmounting 107
- baseline 65
- boot loader
 - reinstalling for Linux after running mbralign 182

C

- C# 165
- capabilities
 - accessing form VSC for VMware vSphere 5
 - interaction 4
 - VSC for VMware vSphere 3
- Citrix 58, 63, 66, 67, 75, 143, 162
- clone 64, 65, 75, 135, 136, 138, 146, 153, 156, 163
 - virtual machines 58
- clone creation engine 135

- clone data 63
- clone operation 63
- cluster 58, 139
- communication
 - troubleshooting 169
- concurrent VMware snapshots 170
- config_mpath
 - setting path selection policy for MSCS guests 175
- configuration 20
 - troubleshooting 170
- configuration files
 - for restore sessions 112, 114
 - generated upon creation of restore sessions 108
 - restore sessions 110
 - used for single file restore 113
- connection broker 63, 64, 66, 83, 142, 162
- Connection brokers panel 63
- connectivity
 - troubleshooting 169
- consistency snapshots 105
- controller 140, 144, 145, 146, 159
 - removing skipped or unmanaged 26
 - supports direct connections to Vserver 24
 - supports vFilter unit, Vserver tunneling 23
- controller name
 - correcting when unknown 25
- controlling what is displayed
 - using Inventory panel 26
- copy/clone offload engine 135
- copyright and trademark information 187
- copyright information 187
- Create Rapid Clones Wizard 66
- createClones 150, 153, 158
- createDatastore 158
- credentials
 - default for storage controllers 21
 - setting for storage controller 22
 - upgrade considerations 14
 - using RBAC 13
- csv 58, 63, 66, 82, 83
- custom user accounts
 - configuring using RBAC 13

D

- data collection
 - log files 169
- Data Collection panel
 - collecting diagnostic data 33
 - gathering information for troubleshooting 51
- data collection programs
 - changing service credentials 34
- datacenter 58, 139
- datastore 82, 84, 135, 136, 139, 140, 141, 144, 145, 146, 153, 156, 158

- datastore management engine 135
- datastores 68, 69, 75, 91, 135
 - cloning 58
 - creating 174
 - creating backups 101
 - mounting 74
 - NFS indirect path 29
 - reclaiming space 65
 - resizing 74
 - restoring a 106
 - searching for backups of 106
 - where to restore a backup 105
- deduplication 67, 75
 - enabling 74
- default credentials
 - for storage controllers 21
- default values
 - overriding 170
- Desktop Studio 67
- destroy datastores 75
- diagnostic data
 - collecting 33
- discovery 20
 - correcting unknown controller name 25
 - enabling for vFiler units 23
- disk partitions 112, 115
 - checking and fixing alignment 35, 36
- Disk tab 112, 115
- Disk.QFullSampleSize 27
- Disk.QFullThreshold 27
- display
 - limiting using Inventory panel 26
- drive letters
 - changing 112, 115

E

- e-mails
 - link to Restore Agent download URL 113
 - restore session confirmations 111
- Element Manager
 - launching 31
- email alerts 100
- email notification error 171
- Emulex FC HBA timeouts 27
- error logs 171
- errors
 - mounting NFS datastore backup fails 171
- ESX 75, 136, 138, 139, 140, 141, 142, 143, 145, 146, 150, 151, 153, 156
- ESX host
 - displaying settings 168
 - multipathing and timeout settings 27
 - restoring virtual disks on VMFS datastores 106
 - setting up 174
 - timeout values 175
- ESX, ESXi host settings 27, 29
- ESXi
 - enable ssh 34
- event logs 171
- expiration times
 - of restore sessions, changing 110

F

- failed consistency snapshots 105
- failover cluster
 - setting path selection policy for 175
- FCP 156, 158
- fileCloneOffload 148, 149
- fileCopyOffload 148, 149
- FilerView
 - launching 31
 - opening 168
- fixed
 - path selection policy 175
- FlexClone 68, 75, 135, 146
- FlexClone license 112, 114
- FQDN 162
- functional alignments 88

G

- General tab
 - changing vCenter Server authentication credentials 115
 - single file restore settings 110
- getFileOpOffloadStatus 149
- GRUB
 - reinstalling for Linux after running mbralign 182
- guest OS
 - installing scripts 37
 - setting timeouts for Linux 176
 - setting timeouts for Solaris 177
 - setting timeouts for Windows 177
 - timeout values 175

H

- Help, online 6
- host
 - multipathing and timeout settings 27
- HTTP 159
- httpd.admin.enable option 23
- HTTPS 159

I

- igroup
 - creating 173
- in-place restore 105
- installation
 - of Restore Agent 112, 114
- installing
 - guest operating system (GOS) scripts 37
 - mbralign 35, 36
 - VSC for VMware vSphere in silent mode 10
 - VSC for VMware vSphere using installation wizard 9
- Inventory panel
 - limiting what is displayed 26
- iSCSI 156, 158
- issues
 - mounting NFS datastore backup fails 171

J

- Job Properties dialog box 103

K

- kaminoprefs file 75, 82

L

- license requirements
 - single file restore feature 112, 114
- limited self-service restore sessions
 - about 109
- limiting what is displayed
 - using Inventory panel 26
- Linux
 - reinstalling GRUB after running mbralign 182
 - setting timeouts for guest OS 176
- linux_gos_timeout-install.iso
 - guest OS tool 176
- Load Configuration window
 - in Restore Agent 112
- location of log files 171
- lock management 6
- Lock management
 - Optimization and Migration 89
- log configuration file 82
 - modifying 82
- log files
 - collecting 169
- log4j 82
- log4j.properties 82
 - log configuration file 82
- logs
 - collecting 169
 - errors 171
 - events 171
 - troubleshooting 171
 - viewing 171
- LUN 67, 68, 75, 156
 - creating 173
 - type 173

M

- MAC address 160
- Managed Object Browser 135
- MBR tools 35, 36
 - enable ssh for ESXi 34
- mbralign
 - checking partition alignment 178
 - fixing partition alignment 181
 - installing 35, 36
 - overview 179
- mbrscan
 - replaced by mbralign 178
- memory requirements
 - Virtual Storage Console 8
- message logs 171
- Monitoring and Host Configuration 19
 - upgrade considerations 14
- mount expiration
 - changing the default time 111, 113

- MSCS
 - setting path selection policy for 175
- multihost configurations
 - backing up 97
 - recovering 97
- multipathing
 - configuring ESX host 27
- MultiStore
 - display differences with vFiler units 32
 - enabling discovery of vFiler units on private networks 23

N

- Net.TcpipHeapMax 27
- Net.TcpipHeapSize 27
- network configuration file 82
- NFS 58, 64, 75, 136, 140, 145, 146, 156, 158
 - indirect path 29
 - reclaiming space 65
- NFS datastore backup
 - mounting a backup fails 171
- NFS paths
 - changing to direct access 30
- NFS VAAI Plug-in
 - installing 30
- NFS.HeartbeatFrequency 27
- NFS.HeartbeatMaxFailures 27
- NFS.HeartbeatTimeout 27
- NFS.MaxVolumes 27
- notices 189
- Notices 189
- NTFS
 - reclaiming space 65
- nvram 75

O

- one-time backup, starting 103
- online Help 6
- Optimization and Migration 91, 93
 - actual alignments 88
 - features 87
 - functional alignments 88
 - important notes 89
 - lock management 6
 - Lock management 89
 - migrating multiple virtual machines 89
 - offline alignments 88
 - online alignments 88
 - SDRS 89
 - VAAI extended copy 89
 - VMware Storage vMotion 89
 - Windows 2008 R2 SP1 89
- out-of-place restore 105
- Overview panel
 - limiting what is displayed 26

P

- parameters
 - ESX, ESXi host 27
 - UNMAP 29

- partition alignment
 - checking with mbralign 178
 - fixing with mbralign 181
 - identifying and fixing 178
- partitions
 - checking and fixing alignment 35, 36
- path selection policy 27
 - setting for MSCS and failover cluster 175
- paths
 - changing to direct NFS paths 30
 - performing online alignments 93
- port groups
 - changing the network connection for 110
 - management of 108
- preferences file 75
- Provisioning and Cloning
 - about 57
 - accessing features 57
 - cloning virtual machines 58
 - lock management 6
 - mounting datastores 74
 - reclaiming space 65
 - upgrade considerations 14
 - using the wsimport tool 164
- Provisioning and Cloning log files 81
- Provisioning and Cloning methods 136
- provisioning storage 173
- PXE 75

Q

- QLogic
 - FC HBA timeouts 27
 - iSCSI HBA IP_ARP_Redirect 27
 - iSCSI HBA timeouts 27

R

- RBAC
 - configuring 13
 - upgrade considerations 14
- RDM 138
- reclaiming space
 - using Provisioning and Cloning 65
- redeploy VMs 150, 153
- registering
 - Virtual Storage Console 11
- Release Notes
 - checking 167
- Remove Controller command 26
- required ports
 - firewall requirements 11
 - Virtual Storage Console 11
- rescan SCSI bus 174
- resizing datastores 74
- resource pool 58
- resources
 - discovering and adding 24
- restore
 - a datastore 106
- Restore Agent
 - changing the download URL 110
 - clearing the configuration 113, 116
 - disk partitions 112, 115

- Restore Agent (*continued*)
 - downloading the application software 113
 - installing 112, 114
 - link to software download 111
 - loading the configuration file 112
 - network connection to Backup and Recovery capability 108
 - post-installation configuration changes 108
 - restoring virtual disk files 112, 115
 - software requirements of 112, 114
- restore operations
 - example of limited self-service 113
 - from VMDKs 106
 - restoring from backups with failed VMware consistency snapshots 105
 - troubleshooting 170
 - types of 105
 - where to restore a backup 105
- Restore panel 105
 - searching for backups 106
- restore sessions
 - Admin Assisted 110
 - changing expiration times 110
 - configuration files 110
 - creating a new session after changing port group settings 108
 - example of a limited self-service restore session 113
 - types of 109
- restore sessions,
 - example of a self-service restore session 111
- Restore wizard
 - restoring a virtual machine or its disk files 106
- restoring virtual data
 - of virtual disk files 108
- resume, suspended backup job 104
- roles
 - configuring with RBAC 13

S

- scanning datastores 91
- scanning with the Optimization and Migration capability 91
- scheduled backup jobs
 - deleting 103
- scheduled backups
 - editing 103
- scripts 100
- scripts, guest operating system (GOS)
 - installing 37
- SCSI bus
 - rescan 174
- SDRS
 - Optimization and Migration 89
- Search field
 - using the 106
- search for a backup 106
- security
 - configuring using RBAC 13
- self-service restore sessions
 - about 109

- service account
 - changing credentials 34
- settings
 - ESX, ESXi host 27, 29
- Setup panel
 - configuring vCenter Server 115
 - single file restore settings 110
- SFR files
 - See* configuration files
 - clearing 113, 116
 - loading with Restore Agent 112
 - show details
 - displaying ESX host settings 168
- single file restore
 - recovery of virtual disk files 108
- single file restore feature
 - license requirements 112, 114
- Single File Restore panel
 - changing the type of restore session 110
 - checking the port group settings 108
 - viewing the restore sessions list 110, 111, 113
- skipped controller
 - removing 26
- SMVI log viewer
 - enabling network adapters 171
- smvi.override file 170
- SnapManager for Virtual Infrastructure server
 - manually changing the IP address 110
- SnapMirror requirements
 - Backup and Recovery 97
- snapshot 75, 136, 138, 153
- Snapshot and SnapRestore requirements
 - Backup and Recovery 97
- snapshot autodelete 68
- snapshots 105
- SOAP 164, 165
- SOAP (Simple Object Access Protocol)
 - programmatic access 135
- software requirements
 - for Restore Agent 112, 114
- Solaris
 - setting timeouts for guest OS 177
- solaris_gos_timeout-install.iso
 - guest OS tool 177
- spanned entities 100
- ssh
 - enabling for ESXi 34
- SSL certificate
 - regenerating 12
- status
 - getting more information for alert 168
- status reason
 - displaying column 168
- storage controller 67, 68, 75
 - removing skipped or unmanaged 26
 - setting default credentials 22
 - using default credentials 21
- storage controller exports file 75
- storage controllers panel 67
- storage resources
 - discovering and adding 24

- storage system
 - discovery 20
 - managing with Element Manager 31
 - managing with FilerView 31
 - managing with System Manager 31
- storage systems
 - configuring using RBAC 13
- supported configurations
 - memory requirements 8
- SUSE Linux
 - reinstalling GRUB after running mbralign 182
- suspend, backup job 104
- suspended backup job, resuming 104
- sysprep answer file 64
- System Manager
 - managing storage controllers 31

T

- the managed object reference of the requested object based on name and type.vCenter server 151
- thick volume 68
- thin provisioned LUN 68
- timeout settings
 - configuring ESX host 27
- timeout values
 - recommended values 175
 - setting for guest OS 175
- timeouts
 - ESX, ESXi host 27, 29
- tools
 - partition alignment 35, 36
 - setting Linux guest OS timeouts 176
 - setting Solaris guest OS timeouts 177
 - setting Windows guest OS timeouts 177
- trademark information 188
- troubleshooting 167
 - checking Release Notes 167
 - collecting diagnostic data 33
 - collecting log files 169
 - email notification error 171
 - gathering information 51
 - mounting NFS datastore backup fails 171
 - partition alignment 178
 - unable to communicate with the server 169
- tunneling
 - supported for vFiler units 99
 - supported for vFiler units, Vservers 23

U

- unable to communicate with the server
 - troubleshooting 169
- uninstalling
 - VSC for VMware vSphere using a command line 17
 - VSC for VMware vSphere using Add/Remove Programs 17
- unknown
 - controller name 25

- unmanaged controller
 - removing 26
- update command 20
- updating
 - resource information 24
- upgrading VSC for VMware vSphere
 - considerations 14
- URL
 - Restore Agent downloads
 - changing 110
- user accounts
 - creating custom users 99
 - creating roles, groups, and users 98
- user accounts, creating roles, groups, and users 99
- user credentials
 - backup job 100
- user name
 - configuring custom with RBAC 13
- UUID 75, 82

V

- VAAI extended copy
 - Optimization and Migration 89
- vApp 58
- vApps 66, 135
 - provisioning and cloning 135
- vCenter 84
- vCenter Inventory 135
- vCenter object 156
- vCenter server 136, 138, 139, 140, 141, 142, 143, 145, 146, 150
- vCenter Server 63, 64, 65, 67, 68
 - registering Virtual Storage Console 11
- vCenter Servers
 - configuring for single file restore 115
 - managing credentials 135
- vCenter Servers, configuration of 99
- vCenter task 136, 138, 140, 141, 146
- vCloud tenants, discovering objects for 135
- vFiler unit
 - display differences with physical storage controllers 32
 - enabling discovery 23
 - tunneling supported 23
- vFiler units
 - discovering on private networks 23
 - tunneling 99
- VI SDK 136
- Virtual Center 136, 138
- Virtual Desktops page 66
- virtual disk files
 - administrator-assisted restoration of 108
 - recovery using Restore Agent 112, 115
- virtual entities 100
- virtual hard disk 156
- virtual hard drive 135, 138
- virtual hard drives 153
- virtual machine 64, 66, 67, 84, 135, 136, 138, 142, 143, 150, 153, 160, 161
 - cloning 58

- virtual machine disk files
 - where to restore a backup 105
- virtual machines
 - creating backups 100
 - migrating in a group 94
 - migrating with Optimization and Migration 89
 - restarting 106
 - restoring 106
 - searching for backups of 106
 - where to restore a backup 105
- virtual networks
 - connection to 108
 - relationship of port groups to network connectivity 108
- Virtual Storage Console 84, 162
 - Backup and Recovery capability 97
 - CLI 116
 - configuration 20
 - firewall port requirements 11
 - installation overview 7
 - launching the VSC CLI 116
 - memory requirements 8
 - registering 11
 - required ports 11
 - supported configurations 8
- Virtual Storage Console for VMware vSphere service
 - changing credentials 34
- VM BIOS file 75
- vmdk 146, 156
- VMDK 58
- VMDK partition alignment
 - checking with mbralign 178
 - fixing with mbralign 181
 - identifying and fixing 178
- VMDKs 106
 - restoring 106
- VMFS 58, 136, 140, 146, 156, 158
- VMFS datastores
 - restoring virtual disks on 106
- VMkernel 75
- VMware 160
- VMware Session ID 153
- VMware snapshots 105
- VMware Storage vMotion
 - Optimization and Migration 89
- VMware vCenter SDK 153
- VMware vCloud
 - APIs for 135
- VMware VI API 153, 156
- VMware VI SDK 135
- VMware View 64, 162
- VMware View Server 58, 63, 142
- vmx 156
- VMX Path 82
- volume 67, 68, 156
 - adding to resources 67
 - creating 173
- volume autogrow 68
- VSC CLI 116
 - launching 116
- VSC for VMware vServer
 - limitations with direct connections to Vservers 24
- VSC for VMware vSphere
 - accessing capabilities 5
- VSC for VMware vSphere (*continued*)
 - architecture 4
 - capabilities 3
 - installing in silent mode 10
 - installing using installation wizard 9
 - interaction between capabilities 4
 - lifecycle management for VMware environments 3
 - lock management 6
 - Monitoring and Host Configuration capability 19
 - Optimization and Migration capability 87
 - plug-in functions 3
 - Provisioning and Cloning capability 57
 - regenerating an SSL certificate 12
 - uninstalling using a command line 17
 - uninstalling using Add/Remove Programs 17
 - upgrade considerations 14
 - upgrading the software 15
- VSC for VMware vSphere Console
 - adding or removing Backup and Recovery capability 16
- Vserver
 - limitations when connecting directly to storage 24
 - supports direct connections to storage 24
 - tunneling supported 23

W

- Windows
 - setting timeouts for guest OS 177
- Windows 2008 R2 SP1
 - Optimization and Migration 89
- Windows failover cluster
 - setting path selection policy for 175
- windows_gos_timeout.iso
 - guest OS tool 177
- workflow 91
- WSDL 164
- wsimport certificate 164

X

- XenDesktop 58, 63, 66, 67, 83, 143, 162

Z

- ZAPI 159



NA 210_05530_A0, Printed in USA

GC52-1349-06

