Veritas InfoScale™ Operations Manager 7.3 Add-ons User's Guide

VERITAS™

Last updated: 2017-06-01

Legal Notice

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

http://www.veritas.com

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

https://www.veritas.com/support

You can manage your Veritas account information at the following URL:

https://my.veritas.com

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

https://sort.veritas.com/documents

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

doc.feedback@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

http://www.veritas.com/community/

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Section 1 VCS Utilities Add-on 7.3 13

Chapter 1 Overview of VCS Utilities - Manage VCS User Passwords Add-on 14

About VCS Utilities - Manage VCS User Passwords Add-on 14
Features of VCS Utilities - Manage VCS User Passwords Add-on 15
Limitations of VCS Utilities - Manage VCS User Passwords Add-on 15
Prerequisites for using VCS Utilities - Manage VCS User Passwords Add-on 16

Chapter 2 Using VCS Utilities Add-on 17

Changing a user's password for non-secure Cluster Server-managed clusters 17
Change a user's password options 18

Section 2 Distribution Manager Add-on 7.3 19

Chapter 3 Introduction to Distribution Manager Add-on 20

About Distribution Manager Add-on 20

Chapter 4 Using Distribution Manager Add-on 21

Creating a customized solution in Veritas InfoScale Operations Manager 21
Sample scripts for customized solutions 23

Section 3 Fabric Insight Add-on 7.3 25

Chapter 5 Overview of Fabric Insight Add-on 7.3 26

About storage area network fabric discovery using Fabric Insight Add-on 26
About setting up switch discovery 27
About Cisco switch discovery 28
About Brocade switch discovery 28
About troubleshooting fabric discovery 29

Chapter 6 Configuring switch discovery using Fabric Insight Add-on 31

Adding a Cisco switch 31
Add switch panel options for selecting vendor and discovery method 32
Add switch panel options for specifying discovery host and other information for Cisco switch discovery 32
Add switch panel options to provide configuration name 33
Adding a Brocade switch 34
Add switch panel options for specifying discovery host and other information for Brocade switch discovery 35

Chapter 7 Managing switch discovery using Fabric Insight Add-on 37

Refreshing switch discovery 37
Editing switch configuration 38
Removing switch discovery 38

Section 4 Patch Installer Add-on 7.3 39

Chapter 8 Introduction to Patch Installer Add-on 40

About Patch Installer Add-on 40

Chapter 9 Using Patch Installer Add-on 42

Prerequisites for deploying SFHA hot fixes 42
Deploying SFHA hot fixes 43
Requirements for scripts to customize SFHA hot fix deployment 44
Adding pre-installation and post-installation scripts to SFHA hot fixes 46

Removing or replacing custom scripts for SFHA hot fixes 47
Viewing information about SFHA hot fix deployment requests 48

Section 5 Storage Insight Add-on 7.3 49

Chapter 10 Performing the deep discovery of enclosures 50

About Storage Insight Add-on 51
About discovery methods and discovered information of storage enclosures 52
About the discovery host 53
About the network-attached storage discovery 54
Adding HITACHI storage enclosures for deep discovery 54
Add Enclosure panel options for vendor and product selection 55
Add Enclosure panel options for enclosure selection 56
Add Enclosure panel options to update the enclosures discovery information 66
Adding EMC Symmetrix storage enclosures for deep discovery 68
Adding IBM XIV storage enclosures for deep discovery 69
Adding NetApp storage enclosures for deep discovery 70
Adding EMC CLARiiON storage enclosures for deep discovery 72
Adding HP EVA storage enclosures for deep discovery 73
Adding IBM System Storage DS enclosures for deep discovery 74
Adding EMC Celerra storage enclosures for deep discovery 75
Adding EMC VNX storage enclosures for deep discovery 76
Adding EMC VPLEX storage enclosures for deep discovery 78
Adding 3PAR storage enclosures for deep discovery 79
Adding IBM SVC storage enclosures for deep discovery 80
Editing the deep discovery configuration for an enclosure 81
Edit Configuration panel options to modify the device discovery 82
Edit Configuration panel options to select an enclosure for deep discovery 93
Removing deep discovery for a storage enclosure 94
Refreshing the enclosures that are configured for deep array discovery 95
Monitoring the usage of thin pools 96
Monitor Thin Pool panel options 97

Chapter 11 Monitoring storage array metering data 99

About performance metering statistics for enclosure 99
Disabling performance metering for an enclosure 101
Enabling performance metering for an enclosure 101
Viewing the performance graphs for an enclosure 102
Viewing the performance graphs for an array port 103
Viewing the performance graphs for an adapter 104

Chapter 12 Managing LUN classifications 106

About LUN classification 106
Creating LUN classifications 107
Modifying LUN classifications 108
Deleting LUN classifications 109
Modifying the order of the LUN classifications 110
Refreshing the LUN classifications 110

Appendix A Enclosure configuration prerequisites 111

HITACHI enclosure configuration prerequisites 111
Physical connection requirements for HITACHI enclosure 112
EMC Symmetrix storage array configuration prerequisites 112
Physical connection requirements for EMC Symmetrix enclosure 112
Device setup requirements for EMC Symmetrix arrays 113
Veritas InfoScale Operations Manager setup requirements for in-band EMC Symmetrix storage arrays 116
Veritas InfoScale Operations Manager setup requirements to discover EMC Symmetrix storage arrays through remote SYMAPI servers 117
IBM XIV enclosure configuration prerequisites 117
Physical connection requirements for IBM XIV enclosure 117
Device setup requirements for IBM XIV enclosure 117
NetApp storage enclosure configuration prerequisites 118
Physical connection requirements for NetApp enclosure discovery 118
Device setup requirements for NetApp enclosure discovery 119
EMC CLARiiON storage enclosures configuration prerequisites 119
Physical connection requirements for EMC CLARiiON enclosure 119
Device setup requirements for EMC CLARiiON enclosure 120
Hewlett-Packard Enterprise Virtual Array (HP EVA) configuration prerequisites 122

Physical connection requirements for HP EVA enclosure 122
Device setup requirements for HP EVA enclosure 123
Verifying CLI functionality for HP EVA enclosure 123
Verifying SSSU CLI communication with HP Command View EVA software 124
Using password file to access Command View EVA software 125
IBM System Storage DS enclosure configuration prerequisites 125
Physical connection requirements for IBM System Storage DS enclosure 125
Device setup requirements for IBM System Storage DS enclosure 126
Using DSCLI Password Security file for System Storage DS enclosure 126
IBM SVC enclosure configuration prerequisites 126
Physical connection requirements for IBM SVC enclosure 127
Using SSH private key file for IBM SVC enclosure 127
EMC Celerra enclosure configuration prerequisites 127
Physical connection requirements for EMC Celerra enclosure 127
Device setup requirements for EMC Celerra enclosure 128
EMC VNX storage enclosure configuration prerequisites 131
Using enclosure credentials for EMC VNX storage array discovery 131
EMC VPLEX storage enclosure configuration prerequisites 132
Physical connection requirements for EMC VPLEX enclosure 132
Device setup requirements for EMC VPLEX enclosure 132
3PAR storage enclosure configuration prerequisites 133

Appendix B Commands used by Management Server for deep discovery of enclosures 134

HITACHI storage enclosure commands 135
EMC Symmetrix storage enclosure commands 135
IBM XIV storage enclosures commands 137
NetApp storage enclosure commands 139
EMC CLARiiON storage enclosure commands 140
HP EVA storage enclosure commands 141
IBM System Storage DS enclosure commands 142
EMC Celerra storage enclosure commands 145
EMC VNX (Block) storage enclosure commands 146
EMC VNX (File) storage enclosure commands 147

EMC VPLEX storage enclosure commands 148
3PAR storage enclosure commands 149
IBM SVC storage enclosure commands 149

Section 6 Storage Insight SDK Add-on 7.3 151

Chapter 13 Overview of Storage Insight SDK Add-on 7.3 152

About Storage Insight SDK Add-on 152
Array objects discovered by Storage Insight plug-in 153
Operational workflow to create Storage Insight plug-in 153
About discovering array information using Storage Insight SDK Add-on 154

Chapter 14 Managing Veritas InfoScale Operations Manager Storage Insight plug-ins 155

About creating Storage Insight plug-in 155
About installing Storage Insight SDK Add-on 157
About discovery script 157
About the enclosure discovery command output 161
About additional scripts 180
About device identifiers 181
Storage Insight Plug-in sample 182
Creating a Storage Insight plug-in 183
Create, edit, and test Storage Insight plug-in panel options 183
Storage Insight Plug-in Vitals panel options 184
Configuration Parameters panel options 185
Enclosure Capabilities panel options 186
Device Identifier panel options 190
Confirmation panel options 191
Editing a Storage Insight plug-in 192
Upload Storage Insight Plug-in panel options 193
Testing a Storage Insight plug-in 193
Upload Storage Insight Plug-in panel options 195

Section 7 Storage Provisioning and Enclosure Migration Add-on 7.3 196

Chapter 15 Provisioning storage 197

About storage provisioning 197
About creating a storage template 198
Creating a storage template using VxFS file systems 199
Create Storage Template – File system parameters panel options 200
Create Storage Template – File system advanced panel options 200
Create Storage Template – Volume parameters panel options 203
Select LUN Characteristics options 204
Create Storage Template – Volume advanced panel options 204
Creating a storage template using NTFS file systems 205
Creating a storage template using volumes 206
Updating a storage template 207
Provisioning storage 208
Uploading storage templates 209
Downloading storage templates 209
Deleting storage templates 210
Locking storage templates 210
Unlocking storage templates 210

Chapter 16 Migrating volumes 212

About volume migration 212
About the Impact Analysis report for volume migration 213
Migrating volumes by host 214
Select LUN Characteristics options 216
Change layout options 216
Migrating volumes by enclosure 217
Migrating volumes by disk group 219
Pausing or resuming a volume migration 221
Deleting a scheduled volume migration 221

Section 8 Veritas HA Plug-in for VMware vSphere Web Client 222

Chapter 17 Introduction to Veritas HA Plug-in for vSphere Web Client 223

About Veritas HA Plug-in for vSphere Web Client 223
Limitations of Veritas HA Plug-in for vSphere Web Client 224

Chapter 18 Installation and uninstallation of Veritas HA Plug-in for vSphere Web Client 225

Before installing the Veritas HA Plug-in for vSphere Web Client 225
Installing Veritas HA Plug-in for vSphere Web Client on the Management Server 226
Before uninstalling the Veritas HA Plug-in for vSphere Web Client 227
Uninstalling the Veritas HA Plug-in for vSphere Web Client 227

Chapter 19 Configurations for Veritas HA Plug-in for vSphere Web Client 228

Registering the HA Plug-in with VMware vCenter Server 228
Unregistering the HA Plug-in from VMware vCenter Server 229
Deploying the HA Plug-in if the Management Server is configured in a high availability environment 230
Adding managed hosts to the Management Server 231
Migrating virtual machines to Veritas InfoScale Operations Manager 232

Section 9 Application Migration Add-on 236

Chapter 20 Introduction to Application Migration Add-on 237

About Application Migration add-on 237
Before installing the Application Migration add-on 238
Installing the Application Migration Add-on 238
Before uninstalling the Application Migration add-on 239
Uninstalling the Application Migration add-on 239

Chapter 21 Creating and managing an application migration plan 240

Supported versions and platforms 241
User privileges 242
Prerequisites for creating an application migration plan 243
VVR Replication: Environment variables used in application migration 245
Creating an application migration plan 246
Understanding user-defined tasks 254
Understanding application migration operations 255
Understanding the Setup Storage operation 255
Understanding the Rehearse operation 257
Understanding the Migrate operation 258
Understanding the cleanup operation 259
Understanding the tasks executed in each operation 259
Validations performed before migration plan execution 268
Executing the application migration plan 270
Editing an application migration plan 272
Deleting application migration plan(s) 273
Exporting application migration plan(s) 273
Importing application migration plan(s) 274
Viewing historical runs 275
Viewing properties of an application migration plan 275
Application migration logs 276

Index 277

Section 1

VCS Utilities Add-on 7.3

- Chapter 1. Overview of VCS Utilities - Manage VCS User Passwords Add-on
- Chapter 2. Using VCS Utilities Add-on

Chapter 1

Overview of VCS Utilities - Manage VCS User Passwords Add-on

This chapter includes the following topics:

- About VCS Utilities - Manage VCS User Passwords Add-on
- Features of VCS Utilities - Manage VCS User Passwords Add-on
- Limitations of VCS Utilities - Manage VCS User Passwords Add-on
- Prerequisites for using VCS Utilities - Manage VCS User Passwords Add-on

About VCS Utilities - Manage VCS User Passwords Add-on

The VCS Utilities - Manage VCS User Passwords Add-on lets Veritas InfoScale Operations Manager administrators change a user's cluster authentication password for one or more clusters. Using the Veritas InfoScale Operations Manager console, you can select clusters, and enter the desired user's name in the **Change Password for the User** panel. You can enter a new cluster authentication password that meets the password complexities requirements for one or more clusters for the selected user.

You can perform this operation only on the non-secure clusters that are managed by Cluster Server. The provision to select multiple users and change their respective cluster authentication passwords is not available.

For information on the Veritas InfoScale Operations Manager versions that the add-on is compatible with, refer to the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

See "Features of VCS Utilities - Manage VCS User Passwords Add-on" on page 15.

See "Limitations of VCS Utilities - Manage VCS User Passwords Add-on" on page 15.

See "Changing a user's password for non-secure Cluster Server-managed clusters" on page 17.

Features of VCS Utilities - Manage VCS User Passwords Add-on

Veritas InfoScale Operations Manager is a comprehensive management platform for cluster environments and provides you with the information about clusters in your data center. However, creating and resetting passwords for multiple users for multiple clusters can become a time consuming and error-prone task. Using the VCS Utilities Add-on, administrators can create and edit a user's password for one, multiple, or all clusters in a single operation. This automation reduces administrative efforts to a great extent.

See "About VCS Utilities - Manage VCS User Passwords Add-on" on page 14.

See "Limitations of VCS Utilities - Manage VCS User Passwords Add-on" on page 15.

Limitations of VCS Utilities - Manage VCS User Passwords Add-on

This section lists the limitations of VCS Utilities - Manage VCS User Passwords Add-on.

- You can change the authentication passwords for only those users who belong to the non-secure clusters that are managed by Cluster Server.
- You can change the cluster authentication password for only one user in a single operation.

See "About VCS Utilities - Manage VCS User Passwords Add-on" on page 14.

See "Features of VCS Utilities - Manage VCS User Passwords Add-on" on page 15.

Prerequisites for using VCS Utilities - Manage VCS User Passwords Add-on

Before you start using VCS Utilities Add-on, ensure that the following pre-requisites are met:

- The version of Veritas InfoScale Operations Manager installed on Management Server must be 4.0, or later.
- The version of Veritas InfoScale Operations Manager installed on the cluster nodes (managed hosts) must be 3.0, or later.
- You are logged on as domain administrator.
- The user for whose account you want to create or edit passwords, must already be a member of the non-secure VCS-managed clusters. The user must be the member of at least one such cluster and may have access to many clusters.
- The clusters must not be in offline or faulted state.

See "About VCS Utilities - Manage VCS User Passwords Add-on" on page 14.

See "Features of VCS Utilities - Manage VCS User Passwords Add-on" on page 15.

See "Limitations of VCS Utilities - Manage VCS User Passwords Add-on" on page 15.

Chapter 2

Using VCS Utilities Add-on

This chapter includes the following topics:

- Changing a user’s password for non-secure Cluster Server-managed clusters
- Change a user’s password options

Changing a user’s password for non-secure Cluster Server-managed clusters

Use VCS Utilities - Manage VCS User Passwords Add-on to change a user's cluster authentication password collectively for multiple clusters. You can select one or more non-secure VCS-managed clusters, and edit a user's password. Ensure that the selected user is already a member of one or more non-secure VCS-managed clusters.

To change a user's password for non-secure Cluster Server-managed clusters

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Do one of the following:
   - To change password for single or multiple clusters:
     - Select **Data Center**, or the Organization. Select the **Clusters** tab.
     - Select one or more clusters and right-click, then select **Change password for cluster user**.
   - To change password for a single cluster:
     - Expand the Organization, or **Uncategorized Clusters** to locate the cluster.
     - Right-click on the required cluster and select **Change password for cluster user**.

3 In the **Change Password for the User** panel, enter the user name and password for the user.

See "Change a user's password options" on page 18.

4 Click **Finish**.

See "About VCS Utilities - Manage VCS User Passwords Add-on" on page 14.

See "Features of VCS Utilities - Manage VCS User Passwords Add-on" on page 15.

See "Limitations of VCS Utilities - Manage VCS User Passwords Add-on" on page 15.

Change a user's password options

The **Change Password for the User** panel displays the selected clusters for which you want to set the password for the desired user. Use this panel to enter the user name and set the new password.

Table 2-1 Change Password for the User options

Fields	Description
Clusters selected	Displays the list of selected clusters. Only the non-secure VCS-managed clusters are supported.
Username	Enter the user name for which you want to set the password.
Password	Enter the new password.
Re-enter Password	Re-enter the new password.

See "Changing a user's password for non-secure Cluster Server-managed clusters" on page 17.

Section 2

Distribution Manager Add-on 7.3

- Chapter 3. Introduction to Distribution Manager Add-on
- Chapter 4. Using Distribution Manager Add-on

Chapter 3

Introduction to Distribution Manager Add-on

This chapter includes the following topics:

- About Distribution Manager Add-on

About Distribution Manager Add-on

Distribution Manager Add-on lets you create a custom solution package, including an install and uninstall routine, to deliver across hosts managed by Veritas InfoScale Operations Manager. You can then deploy the custom solution using the Management Server console in the same way that you deploy standard Veritas InfoScale Operations Manager solutions.

For information on the Veritas InfoScale Operations Manager Management Server versions that the add-on is compatible with, refer to the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

See "Creating a customized solution in Veritas InfoScale Operations Manager" on page 21.

See "Sample scripts for customized solutions " on page 23.

Chapter 4

Using Distribution Manager Add-on

This chapter includes the following topics:

- Creating a customized solution in Veritas InfoScale Operations Manager
- Sample scripts for customized solutions

Creating a customized solution in Veritas InfoScale Operations Manager

In Veritas InfoScale Operations Manager, you can create a customized solution that you can deploy on one or more managed hosts and Management Server. You can use the **Create New Solution** wizard panel to create a customized solution.

Before using the wizard, you must create a setup and unsetup script that you can browse to from the wizard.

The setup script must be named as follows:

```
setup.pl
```

The unsetup script must be named as follows:

```
unsetup.pl
```

See the following topic for script samples.

See “Sample scripts for customized solutions ” on page 23.

To create a customized solution

1. In the Home page on the Management Server console, click **Settings**.
2. Click the **Settings** tab and click **Deployment**.

3 Click **Create New Solution**.

4 In the **Create New Solution** wizard panel, enter the relevant information, as follows:

Name	Enter a name for the customized solution that you create.
Version	Enter a version for the customized solution that you create.
Description	Enter any additional information on the customized solution that you create.
MH Component	Select this check box if you want to deploy the solution on managed hosts.
MS Component	Select this check box if you want to deploy the solution on Management Server.
MH/MS Versions	Select the versions of either the managed hosts, Management Server, or both, on which you want to deploy the solution.
Upload platform specific files:	
Type	Select the platform on which you want to deploy the customized solution that you create. You can select from the following options: ■ Linux ■ SunOS ■ Windows ■ AIX ■ HP-UX
Cross Platform	Select this check box if you want to deploy the solution across platforms. If you select this check box, you cannot select a specific platform in the **Type** field.
Setup Script	Browse to select the setup script for the solution that you create. Click **Upload** to upload the setup script to Veritas InfoScale Operations Manager.

UnSetup Script	Browse to select the unsetup script for the solution that you create. Click **Upload** to upload the unsetup script to Veritas InfoScale Operations Manager.
Additional Files	Browse to select any additional files to include in the solution that you create. Click **Upload** to upload the additional files to Veritas InfoScale Operations Manager.

5 Click **Finish** to create a solution.

See "About Distribution Manager Add-on" on page 20.

Sample scripts for customized solutions

To create customized solutions for distribution using Distribution Manager Add-on, you must upload a setup script and unsetup script for installing and uninstalling the solution. The following are examples of these scripts.

Sample setup script:

```
#!/opt/VRTSsfmh/bin/perl
BEGIN { @INC = ("/opt/VRTSsfmh/lib/modules"); }
# Stage area for the distribution is defined as $stage_dir.
# All files in the distribution will be under $stage_dir
my $stage_dir = $ENV{'Store'};
# Code for installing distribution goes here
# For example
# my $out = `rpm -i $stage_dir/mypackage-1.0.rpm`;
# if($? == 0)
# {
#    exit(0); #Success
# }
# else
# {
#     exit(1); #Failed
# }
```

Sample unsetup script:

```
#!/opt/VRTSsfmh/bin/perl
BEGIN { @INC = ("/opt/VRTSsfmh/lib/modules"); }
# Stage area for the distribution is defined as $stage_dir.
```

```
# All files in the distribution will be under $stage_dir
my $stage_dir = $ENV{'Store'};
# Code for un-installing distribution goes here
# For example
# my $out = `rpm -e mypackage-1.0`;
# if($? == 0)
# {
#    exit(0); #Success
# }
# else
# {
#     exit(1); #Failed
# }
```

Section 3

Fabric Insight Add-on 7.3

- Chapter 5. Overview of Fabric Insight Add-on 7.3
- Chapter 6. Configuring switch discovery using Fabric Insight Add-on
- Chapter 7. Managing switch discovery using Fabric Insight Add-on

Chapter 5

Overview of Fabric Insight Add-on 7.3

This chapter includes the following topics:

- About storage area network fabric discovery using Fabric Insight Add-on
- About setting up switch discovery
- About Cisco switch discovery
- About Brocade switch discovery
- About troubleshooting fabric discovery

About storage area network fabric discovery using Fabric Insight Add-on

You can use Fabric Insight Add-on to discover storage area network (SAN) fabric switches that are configured in the data center.

The Fabric Insight Add-on supports discovery for the following SAN switch vendors:

- Cisco: Using Simple Network Management Protocol (SNMP)
- Brocade: Using HTTP communication or ODBC connection to Brocade Network Advisor (BNA)

Note: You cannot discover an individual switch in the fabric; Fabric Insight Add-on discovers the entire fabric.

Fabric Insight Add-on is supported on the following operating systems:

- 64-bit Linux
- 64-bit Windows

You must install Fabric Insight Add-on on the Management Server and optionally on one or more managed hosts. Any host (Management Server or managed host) where Fabric Insight Add-on is installed, can be used as the discovery host for the fabric discovery. However, it is recommended that you install Fabric Insight Add-on on the managed host and use that managed host for the fabric discovery. The Management Server and managed host must be version 6.0, or later.

Note: The switch discovery happens once every six hours, and Veritas InfoScale Operations Manager database is updated accordingly.

See "About setting up switch discovery" on page 27.

See "About Cisco switch discovery" on page 28.

See "About Brocade switch discovery" on page 28.

About setting up switch discovery

A typical switch discovery configuration involves the following steps:

- Install Fabric Insight Add-on on the Management Server and optionally on one or more managed hosts.
- Configure switch discovery for a fabric.
  - Select the switch vendor and discovery method.
  - Specify the discovery host. The discovery host is used to discover switches.
  - Specify the seed switch. You can designate any one switch in a fabric as the seed switch. A seed switch works as an entry point for the discovery of the other switches in the fabric. A seed switch does not necessarily need the latest firmware on it.
  - Provide the user credentials and other vendor-specific information for discovering the switches.
    For Cisco switches, SNMP communication is used. For Brocade switches, HTTP communication or Brocade Network Advisor (BNA) can be used. For successful discovery, all switches in the fabric must be accessible using the same credentials.
- View the discovered properties of the switch in the **Storage** perspective of the Management Server console. For example, information about switches, switch ports, fabric, zones, zone members, initiator, targets, and other properties.

See “About storage area network fabric discovery using Fabric Insight Add-on” on page 26.

About Cisco switch discovery

Using SNMP communication, the following properties of Cisco switches are discovered:

- VSANs
- Switches (including switches in NPV mode)
- Switch ports
- Active zones and their zone members
- Connectivity information between switches and HBA or storage array ports (FC or FCoE including NPIV)

Cisco switch discovery prerequisites are as follows:

- Supported switch operating systems versions: Cisco SAN OS 3.1 (model DS-C9509 and DS-C9124-K9) and Cisco NX-OS 5.1 (model N5K-C5548UP)
- All switches in the fabric must have the same user credentials.
- Cisco switches in NPV mode cannot be configured as seed switches.
- Discovery of mixed-vendor fabrics is not supported.
- Ensure that the UDP port 161 is open between the discovery host and Cisco switch.

The supported SNMP protocol versions are SNMP v1, v2c, and v3. However, when you select **SNMP v1/v2c** while configuring the switch discovery, SNMP v2c is used for the discovery.

See “About storage area network fabric discovery using Fabric Insight Add-on” on page 26.

About Brocade switch discovery

Brocade switches are discovered using the following methods:

- Using HTTP communication
- Using ODBC connection to Brocade Network Advisor (BNA)

The following information is discovered about the Brocade switches:

- Switches, multi-protocol routers, and the access gateways

- Active zones and their zone members
- Virtual switches
- Inter-switch and NPV/NPIV links between the switches
- Switch to multi-protocol router connections

Brocade switch discovery prerequisites are as follows:

- Principal switch requirement: Fabric OS version 6.3.x, or later
- Subordinate switch requirements: Fabric OS version 5.3.2, or later
- Seed Switch: Fabric OS version 5.3.2, or later
- Minimum guest-level access is required for the discovery.
- All switches in the fabric must have the same user credentials for a successful switch discovery. The BNA credentials are the BNA database credentials; not the BNA application login credentials.
- For the discovery using HTTP communication, ensure that HTTP port 80 is open. For the discovery using Brocade Network Advisor (BNA), postgress port 5432 should be open.
- For BNA based discovery, the BNA version should be 11.x, or later. For HTTP based discovery, the switch firmware version should be 5.3.x, or later.

Note: With some versions of BNA (for example version 12.x and later), additional steps are required on the BNA server to enable the Veritas InfoScale Operations Manager discovery host to discover the switches. Change the BNA database client authentication file (`pg_hba.conf`) to allow the discovery host to connect to the BNA database. The file is located at `C:\Program Files\Network Advisor 12.0.0\data\databases\pg_hba.conf`.

See "About storage area network fabric discovery using Fabric Insight Add-on" on page 26.

About troubleshooting fabric discovery

The fabric discovery may fail due to the following reasons:

- For Cisco:
  - If some of the switches are not reachable or incorrect user credentials are provided for the switches.

- If the seed switch is not reachable, or invalid user credentials are provided for the seed switch.

- For Brocade:
  - If the principal switch is not reachable, or invalid user credentials are provided.
  - If the seed switch is not reachable, or invalid user credentials are provided for the seed switch.
  - BNA services are not running. You can use Network Advisor Server Management Console to verify if all the BNA services are running.
  - BNA server is not reachable via IP from the Veritas InfoScale Operations Manager discovery host. You need to check the firewall and other network configurations.
  - BNA database credentials are not valid.
  - BNA database client authentication file `pg_hba.conf` is not changed (applicable to BNA versions 12 and later).
  - The maximum number of concurrent connections to the BNA database has been reached.

See "About setting up switch discovery" on page 27.

See "About storage area network fabric discovery using Fabric Insight Add-on" on page 26.

Chapter 6

Configuring switch discovery using Fabric Insight Add-on

This chapter includes the following topics:

- Adding a Cisco switch
- Adding a Brocade switch

Adding a Cisco switch

You can use Fabric Insight Add-on to add a Cisco switch to the Management Server.

To add a Cisco switch using Fabric Insight Add-on

1. In the Home page on the Management Server console, click **Settings.**
2. Click **Device**.
3. Do one of the following:
   - Right click **Data Center**, and select **Add Switch**.
   - Expand **Data Center**, right-click on the **Switches** node, and select **Add Switch**.
4. In the **Add Switch** panel, select the vendor. Click **Next**.

   See “Add switch panel options for selecting vendor and discovery method” on page 32.

5 Provide the information on discovery host, seed switch, SNMP version, and user credentials. Click **Next**.

See “Add switch panel options for specifying discovery host and other information for Cisco switch discovery” on page 32.

6 Enter a name for the configuration. Click **Finish**.

See “Add switch panel options to provide configuration name” on page 33.

See “About storage area network fabric discovery using Fabric Insight Add-on” on page 26.

See “About setting up switch discovery” on page 27.

Add switch panel options for selecting vendor and discovery method

Use this wizard panel to select the switch vendor and the discovery method.

Table 6-1 Add switch panel options for selecting vendor and discovery method

Field	Description
Switch Vendor	Select the switch vendor from the drop-down list.
Discovery Method	Select the discovery method. For example, SNMP communication for Cisco switches.

See “Adding a Cisco switch ” on page 31.

See “Adding a Brocade switch ” on page 34.

Add switch panel options for specifying discovery host and other information for Cisco switch discovery

Use this wizard panel to specify the discovery host, seed switch, and other parameters that are required for the discovery of Cisco switches.

Table 6-2 Add switch panel options for specifying discovery host and other information for Cisco switch discovery

Field	Description
Switch Vendor	Displays the switch vendor.
Discovery Method	Displays the discovery method for the discovery of the Cisco switch.

Table 6-2 Add switch panel options for specifying discovery host and other information for Cisco switch discovery *(continued)*

Field	Description
Discovery Host	Select the discovery host.
Seed Switch Name/IP	Enter the name or the IP address of the seed switch that is required for the discovery. For Cisco, a switch in the N Port Virtualization (NPV) mode should not be configured as a seed switch.
SNMP Version	Select the SNMP version to be used for the SNMP communication.
SNMPv3	The options in this section are available only when you select **SNMPv3** in the **SNMP Version** field. User Name: Enter the name of the SNMP user. The user must have at least network-operator role. Security Level: Select the appropriate security level from the drop-down list. The available levels are **NoAuthNoPriv**, **AuthNoPriv**, and **AuthPriv**. Authentication Protocol: Select the authentication protocol. It is applicable to **AuthNoPriv** and **AuthPriv** security levels. Authentication Password: Enter the authentication password for the user. It is applicable to **AuthNoPriv**, and **AuthPriv** security levels. Privacy Protocol: Select the privacy protocol. It is applicable to **AuthPriv** security level. Privacy Password: Enter the privacy password for the SNMP user. It is applicable to **AuthPriv** security level.
SNMPv1/2c	The option in this section is available only when you select **SNMPv1/2c** in the **SNMP Version** field. Community String: Enter the SNMP community string. The community can be read-only or read-write.

See “Adding a Cisco switch ” on page 31.

Add switch panel options to provide configuration name

Use this wizard panel to review the information about the switches that will be discovered.

Table 6-3 Add switch panel options to provide configuration name

Field	Description
Configuration Name	Enter a name for the switch configuration.
Switches to be configured for discovery: This section displays the details of the switches that meet the prerequisites and will be discovered in the console. For example, the display name of the switch, its IP address, discovery state, and the World Wide Name (WWN).	

See "Adding a Cisco switch " on page 31.

See "Adding a Brocade switch " on page 34.

Adding a Brocade switch

You can use Fabric Insight Add-on to add a Brocade switch to the Management Server.

To add a Brocade switch using Fabric Insight Add-on

1. In the Home page on the Management Server console, click **Settings.**
2. Click **Device**.
3. Do one of the following:
   - Right click **Data Center**, and select **Add Switch**.
   - Expand **Data Center**, right-click on the **Switches** node, and select **Add Switch**.
4. In the **Add Switch** panel, select the vendor. Click **Next**.

   See "Add switch panel options for selecting vendor and discovery method" on page 32.
5. Provide the information on discovery host, seed switch, and user credentials. Click **Next**.

   See "Add switch panel options for specifying discovery host and other information for Brocade switch discovery" on page 35.
6. Enter a name for the configuration. Click **Finish**.

   See "Add switch panel options to provide configuration name" on page 33.

See "About storage area network fabric discovery using Fabric Insight Add-on" on page 26.

See "About setting up switch discovery" on page 27.

Add switch panel options for specifying discovery host and other information for Brocade switch discovery

Use this wizard panel to provide the discovery host, seed switch, and other parameters that are required for the discovery of Brocade switches.

If Brocade switch discovery happens through HTTP communication, you need to enter the parameters as explained in the following table:

Table 6-4 Add switch panel options for specifying discovery host and other information for Brocade switch discovery using HTTP communication

Field	Description
Switch Vendor	Displays the switch vendor.
Discovery Method	Displays the discovery method as **Using HTTP Communication**.
Discovery Host	Select the discovery host.
Seed Switch Name/IP	Enter the name or the IP address of the seed switch that is used for the switch discovery.
User Name	Enter the user name to access the discovery host.
Password	Enter the required password to access the discovery host.

If Brocade Network Adviser (using ODBC connection) is used to discover Brocade switch, you need to enter the parameters as explained in the following table:

Table 6-5 Add switch panel options for specifying discovery host and Brocade Network Advisor details for Brocade switch discovery using ODBC connection

Field	Description
Switch Vendor	Displays the switch vendor.
Discovery Method	Displays the discovery method as **Using ODBC connection to Brocade Network Advisor**.
Discovery Host	Select the discovery host that is required for the discovery of Brocade switch using Brocade Network Advisor.
BNA Name/IP	Enter the name or the IP address of the Brocade Network Advisor Server.

Table 6-5 Add switch panel options for specifying discovery host and Brocade Network Advisor details for Brocade switch discovery using ODBC connection *(continued)*

Field	Description
User Name	Enter the user name for Brocade Network Advisor database. The default user name is **dcmuser**. **Note:** The used BNA credentials are the BNA database credentials; not the BNA application login credentials.
Password	Enter the required password to access the Brocade Network Advisor database. The default password is **password**.

See “Adding a Brocade switch ” on page 34.

Chapter 7

Managing switch discovery using Fabric Insight Add-on

This chapter includes the following topics:

- Refreshing switch discovery
- Editing switch configuration
- Removing switch discovery

Refreshing switch discovery

You can use Fabric Insight Add-on to refresh the discovery of SAN switches that are added in the data center.

To refresh the switch discovery using Fabric Insight Add-on

1. In the Home page on the Management Server console, click **Settings.**
2. Click **Device**.
3. Expand **Data Center**, and select the **Switches** node.
4. Under the **Switch Configurations** tab, select the required configuration.
5. Right-click and select **Refresh Configuration**.
6. On the **Refresh Configuration** panel, click **Yes** to confirm.

See "About storage area network fabric discovery using Fabric Insight Add-on" on page 26.

See "About setting up switch discovery" on page 27.

Editing switch configuration

You can use Fabric Insight Add-on to edit the parameter related to the discovery of configured SAN switches.

To edit the switch configuration using Fabric Insight Add-on

1. In the Home page on the Management Server console, click **Settings.**
2. Click **Device**.
3. Expand **Data Center**, and select the **Switches** node.
4. Under the **Switch Configurations** tab, select the required configuration.
5. Right-click and select **Edit Configuration**.
6. Edit the information on discovery host, seed switch, user credentials, and other vendor-specific parameters, and click **Next**.
7. Review the switch-related information, and click **Finish**.

See "About storage area network fabric discovery using Fabric Insight Add-on" on page 26.

See "About setting up switch discovery" on page 27.

Removing switch discovery

You can use Fabric Insight Add-on to remove the discovery of SAN switches that are configured in the data center.

To remove the switch discovery using Fabric Insight Add-on

1. In the Home page on the Management Server console, click **Settings.**
2. Click **Device**.
3. Expand **Data Center**, and select the **Switches** node.
4. Under the **Switch Configurations** tab, select the required configuration.
5. Right-click and select **Remove Configuration**.
6. On the **Remove Configuration** panel, click **Yes** to confirm.

See "About storage area network fabric discovery using Fabric Insight Add-on" on page 26.

See "About setting up switch discovery" on page 27.

Section 4

Patch Installer Add-on 7.3

- Chapter 8. Introduction to Patch Installer Add-on
- Chapter 9. Using Patch Installer Add-on

Chapter 8

Introduction to Patch Installer Add-on

This chapter includes the following topics:

- About Patch Installer Add-on

About Patch Installer Add-on

Veritas InfoScale Operations Manager Management Server can provide information about updates available for Storage Foundation High Availability (SFHA) product versions and platforms that are managed by Veritas InfoScale Operations Manager.

Update information is retrieved from Veritas Services Operations Readiness Tools (SORT). SORT is a website that helps you manage your Veritas products.

You can use the Management Server console to download SFHA updates. Patch Installer Add-on adds the capability of using the Management Server console to deploy and install supported SFHA hot fixes on the managed hosts that require them. Supported hot fixes are those that have been configured as deployable by Veritas InfoScale Operations Manager. The Management Server console shows which hot fixes are deployable.

Patch Installer Add-on is installed on the Management Server only.

You can also use Management Server console to customize an SFHA hot fix by adding custom perl scripts for pre-installation and post-installation tasks. For example, you could add scripts to stop and then restart your applications.

See "Adding pre-installation and post-installation scripts to SFHA hot fixes" on page 46.

More information is available on prerequisites for download and installation using Patch Installer Add-on.

See “Prerequisites for deploying SFHA hot fixes” on page 42.

For information on using Patch Installer Add-on to download and install SFHA hot fixes, see the following topic:

See “Deploying SFHA hot fixes” on page 43.

For more information on SORT, go to:

https://sort.veritas.com

Chapter 9

Using Patch Installer Add-on

This chapter includes the following topics:

- Prerequisites for deploying SFHA hot fixes
- Deploying SFHA hot fixes
- Requirements for scripts to customize SFHA hot fix deployment
- Adding pre-installation and post-installation scripts to SFHA hot fixes
- Removing or replacing custom scripts for SFHA hot fixes
- Viewing information about SFHA hot fix deployment requests

Prerequisites for deploying SFHA hot fixes

In the Management Server console, you can use Patch Installer Add-on to download and install hot fixes for Storage Foundation High Availability (SFHA) products on managed hosts version 6.0 or later.

Supported SFHA hot fixes are shown in the console as **VOM deployable**.

For information about which SFHA updates are supported for installation by Patch Installer Add-on, see the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List* (HSCL).

Downloading hot fixes requires connectivity to the Veritas Services Operations Readiness Tools (SORT) website. If no connection is available to SORT from the Management Server, you can configure a proxy server from the console.

For information on configuring proxy settings, see the *Veritas InfoScale Operations Manager Management Server Installation and Configuration Guide*.

You can also modify the settings for retrieving hot fix information from SORT.

For information on modifying SORT download settings, see the *Veritas InfoScale Operations Manager Management Server User Guide*.

Some hot fixes require a system reboot. During hot fix deployment, you can specify whether or not the system is automatically rebooted after the hot fix is installed. Patch Installer Add-on will detect whether a reboot is required.

You can customize an SFHA hot fix by adding custom perl scripts to handle pre-installation and post-installation tasks. For example, you could add scripts to stop and then restart your applications. You customize the hot fix before deploying it.

See "Adding pre-installation and post-installation scripts to SFHA hot fixes" on page 46.

Deploying SFHA hot fixes

In the Management Server console, you can use Patch Installer Add-on to install supported hot fixes for Storage Foundation High Availability (SFHA) products. Supported SFHA hot fixes are shown in the console as **VOM deployable**.

Note: Patch Installer Add-on is not required to deploy Veritas InfoScale Operations Manager hot fixes.

Some hot fixes require a system reboot. During hot fix deployment, you can specify whether or not the system is automatically rebooted after the hot fix is installed.

Before you begin, review the prerequisites.

See "Prerequisites for deploying SFHA hot fixes" on page 42.

If you want to customize a hot fix by adding custom perl scripts, you should perform that task before deployment.

See "Adding pre-installation and post-installation scripts to SFHA hot fixes" on page 46.

To deploy SFHA hot fixes

1. In the Home page on the Management Server console, click **Settings**.
2. Click **Deployment**.
3. Locate the hot fix as follows:
   - To view hot fixes for all SFHA products, click **Hot Fixes** in the tree and click the **Hot Fixes** tab.

- To view hot fixes for one product only, expand **Hot Fixes** in the tree and click one of the products, then click the **Hot Fixes** tab.
- To filter the list on the **Hot Fixes** tab to show only VOM deployable hot fixes, select the **VOM Deployable** check box.

4. On the **Hot Fixes** tab, select the hot fix to be deployed.
5. Right-click and select **Install**.
6. If the **Download Hot Fix** panel is displayed in the Install wizard, select whether you are downloading from SORT or using a local copy and click **Next**.
7. On the **Select hosts** panel, select the hosts on which the hot fix is to be installed and click **Next**.
8. Optionally, select **Automatically reboot systems after installation is complete**. Click **OK**.
9. After the install request is sent, you can exit the wizard and track the status using the **Recent Tasks** list. To view more details, such as the output of a custom script, double-click on a task. If multiple hosts are being installed, you can expand the task and view details on each host. You can also view information on the **Applicable Hosts** tab.

Requirements for scripts to customize SFHA hot fix deployment

You can create custom perl scripts to perform tasks that are required before and after installation of a Storage Foundation High Availability (SFHA) hot fix by Patch Installer Add-on.

The perl scripts must meet the following requirements:

- The script Exit must return 0 for success.
- The script Exit must return a value other than 0 for failure.

To help you verify that scripts executed successfully, anything printed on STDERR or STDOUT from the script is displayed in the task log details when you install the customized hot fix.

Note: Post-installation scripts are intended for use in cases where hot fixes do not require a reboot. In cases where hot fixes require a reboot, the post-installation script is not executed.

The following are examples showing the content required for using the scripts to customize deployment of Storage Foundation High Availability (SFHA) hot fixes.

Example for pre-installation script:

```
# stop relevant applications before the core services are stopped and
  patch is installed

# exit the script with proper exit code and message.
# if pre setup fails HF/CP installation will be aborted.

$ret = 0;
if ($ret) # zero exit code indicates success. non-zero indicates error.
{

   my $msg = "some error message indicating what failed";

   print STDERR "Pre setup failed: $msg";

   exit(1);
}
else
{
   print STDOUT "Pre setup successfull";
   exit(0);
}
```

Example for post-installation script:

```
## Post setup script. executed after HF/CP is installed but
 before system reboots (if reboot is required)



# start relevant applications which were stopped in the pre-setup
 script.

$ret = 0;
if ($ret) # zero exit code indicates success. non-zero indicates error.
{

   my $msg = "some error message indicating what failed";

   print STDERR "Post setup failed: $msg";
```

```
    exit(1);
}
else
{
    print STDOUT "Post setup successfull";
    exit(0);
}
```

See "Adding pre-installation and post-installation scripts to SFHA hot fixes" on page 46.

Adding pre-installation and post-installation scripts to SFHA hot fixes

You can create custom perl scripts to perform tasks that are required before and after installation of a Storage Foundation High Availability (SFHA) hot fix by Patch Installer Add-on. For example, you could create scripts to stop and restart your applications.

Note: Post-installation scripts are intended for use in cases where hot fixes do not require a reboot. In cases where hot fixes require a reboot, the post-installation script is not executed.

In the Management Server console, you can use Patch Installer Add-on to add the custom scripts to the SFHA hot fix. You add the scripts to customize the hot fix before deploying the hot fix to the hosts.

See "Requirements for scripts to customize SFHA hot fix deployment" on page 44.

To add pre-installation and post-installation scripts to SFHA hot fixes

1. In the Home page on the Management Server console, click **Settings**.
2. Click **Deployment**.
3. Locate the hot fix for customization as follows:
   - To view hot fixes for all SFHA products, click **Hot Fixes** in the tree and click the **Hot Fixes** tab.
   - To view hot fixes for one product only, expand **Hot Fixes** in the tree and click one of the products, then click the **Hot Fixes** tab.
4. On the **Hot Fixes** tab, select a hot fix to be customized.
5. Right-click and select **Customize**.

6 If you have not yet downloaded the hot fix, the wizard panel will prompt you to select whether you are downloading from SORT or using a local copy. Select the appropriate option and click **Next**.

7 If downloading, wait for the download to complete and click **Next**.

8 Upload a pre-installation and/or post-installation script. Browse to the file that you want to add and click **Upload**.

The uploaded file is listed in the table. You can click on the page icon in the table row to view the content.

9 Click **Finish**. When the operation shows that it is completed successfully, click **OK**.

In the **Properties** pane for the selected hot fix, the **User Customized** property will display as **Yes**.

See "Removing or replacing custom scripts for SFHA hot fixes" on page 47.

Removing or replacing custom scripts for SFHA hot fixes

In the Management Server console, you can use Patch Installer Add-on to remove or replace custom pre-installation and post-installation scripts that were added to an SFHA hot fix.

To replace a script, you first remove the script file that you want to replace. Then add the script file that you want to replace it with.

To remove or replace custom scripts for SFHA hot fixes

1 In the Home page on the Management Server console, click **Settings**.

2 Click **Deployment**.

3 Locate the hot fix as follows:

- To view hot fixes for all SFHA products, click **Hot Fixes** in the tree and click the **Hot Fixes** tab.
- To view hot fixes for one product only, expand **Hot Fixes** in the tree and click one of the products, then click the **Hot Fixes** tab.

4 Right-click the hot fix and select **Customize**.

5 On the **Customize** panel, click the remove icon in the table row for the file that you want to remove.

6 If you want to replace a removed script, browse to the replacement script file and upload it.

7 Click **Finish**.

See "Adding pre-installation and post-installation scripts to SFHA hot fixes" on page 46.

Viewing information about SFHA hot fix deployment requests

You can view information about previous requests to deploy Storage Foundation High Availability (SFHA) hot fixes using Patch Installer Add-on. The information includes time, status, number of hosts successfully updated, number of hosts where the update failed, and the names of the hosts.

To view information about SFHA hot fix deployment requests

1 In the Home page on the Management Server console, click **Settings**.

2 Click **Deployment**.

3 In the tree, under **Repository**, click **Hot Fixes**.

4 Click the **Requests** tab.

See "Deploying SFHA hot fixes" on page 43.

Section 5

Storage Insight Add-on 7.3

- Chapter 10. Performing the deep discovery of enclosures
- Chapter 11. Monitoring storage array metering data
- Chapter 12. Managing LUN classifications
- Appendix A. Enclosure configuration prerequisites
- Appendix B. Commands used by Management Server for deep discovery of enclosures

Chapter 10

Performing the deep discovery of enclosures

This chapter includes the following topics:

- About Storage Insight Add-on
- About the discovery host
- About the network-attached storage discovery
- Adding HITACHI storage enclosures for deep discovery
- Adding EMC Symmetrix storage enclosures for deep discovery
- Adding IBM XIV storage enclosures for deep discovery
- Adding NetApp storage enclosures for deep discovery
- Adding EMC CLARiiON storage enclosures for deep discovery
- Adding HP EVA storage enclosures for deep discovery
- Adding IBM System Storage DS enclosures for deep discovery
- Adding EMC Celerra storage enclosures for deep discovery
- Adding EMC VNX storage enclosures for deep discovery
- Adding EMC VPLEX storage enclosures for deep discovery
- Adding 3PAR storage enclosures for deep discovery
- Adding IBM SVC storage enclosures for deep discovery
- Editing the deep discovery configuration for an enclosure

- Removing deep discovery for a storage enclosure
- Refreshing the enclosures that are configured for deep array discovery
- Monitoring the usage of thin pools

About Storage Insight Add-on

Storage Insight Add-on, with its deep array discovery capabilities, provides you the detailed information about the storage enclosures in the data center. Storage Insight Add-on lets you view the detailed storage information on the enclosure to the level of adapters and disks.

For information on the Veritas InfoScale Operations Manager Management Server versions that the add-on is compatible with, refer to the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

After you configure deep discovery for the enclosures, the discovery happens once every six hours, and Veritas InfoScale Operations Manager database is updated accordingly.

Note: To perform deep array discovery of enclosures, you require the access to the storage views and operations.

The Storage Insight Add-on also lets you define the classifications for the LUNs in the data center. You can classify the LUNs based on one or more parameters. For example, enclosure name, vendor, RAID level of a LUN, and the LUN type. The Storage Insight Add-on also lets you perform the maintenance of the Dynamic Multi-Pathing paths in array ports and adapters.

See “About the discovery host” on page 53.

See “About discovery methods and discovered information of storage enclosures” on page 52.

See “About the network-attached storage discovery” on page 54.

See “About performance metering statistics for enclosure” on page 99.

See “About LUN classification” on page 106.

See “Removing deep discovery for a storage enclosure” on page 94.

About discovery methods and discovered information of storage enclosures

Veritas InfoScale Operations Manager Storage Insight Add-on, with its deep discovery capabilities lets you discover the following enclosures.

Table 10-1 Enclosure discovery methods and the discovered information

Enclosure	Discovery method	Discovered information
HITACHI	HiCommand server	LDEVs (logical devices), PDEVs (physical devices), array groups, replications, adapters, host associations, and thin pools.
IBM XIV	XCLI utility	Pool volumes, physical disks, storage pools, replications, thin pools, array ports, adapters, and host associations.
EMC Symmetrix	SymCLI utility	Physical disks, devices, replications, thin pools, array ports, adapters, host associations, Fully Automated Storage Tiering (FAST) technology managed storage groups (VMAX only), and Fully Automated Storage Tiering for Virtual Pools (FAST VP).
EMC CLARiiON	NaviSphere Secure CLI (NaviSecCLI) utility	Thin pools, thin LUNs, physical disks, RAID groups, replications, array ports, adapters, host associations, storage processors, and LUNs.
NetApp	NetApp Data ONTAP SDK	Physical disks, aggregates, array volumes, array ports, adapters, host associations, replications, flexi / traditional volumes, and LUNs. NAS discovery: Qtree, Share, Quota, vFiler, and consumer storage pool
HP Enterprise Virtual Array (EVA)	SSSU CLI utility. It gets the information from the Command View Server.	Virtual disks, physical disks, array disk groups, array ports, adapters, host associations, and replications.
EMC Celerra	Celerra XML API and CLI	Disk volumes, storage pools, array ports, adapters, host associations, array volumes (Meta, Strip, Slice and Pool Volumes), NAS file systems, Shares, and Data Movers.

Table 10-1 Enclosure discovery methods and the discovered information *(continued)*

Enclosure	Discovery method	Discovered information
IBM System Storage DS	DS Command-line interface (DSCLI)	Logical volumes, disk drive modules (DDMs), extent pools, replications, array ports, adapters, host associations, and ranks.
EMC VNX	Using XML API and CLI	LUNs, physical disks, RAID Groups, replications, thin pools, disk volumes, storage pools, array ports, adapters, host associations, array volumes, Datamovers, and NAS objects (for example, Shares). You can discover all three types of VNX configurations - File, Block, and Unified.
EMC VPLEX	Using HTTPS access to the VPLEX server and SSH communication	Virtual volumes, storage volumes, RAID Groups, array ports, adapters, host associations, and replications. Internal nodes, local and distributed devices, local and distributed virtual volumes.
3PAR	Using SSH communication	Virtual volumes, physical disks, replications, array ports, adapters, host associations, Common Provisioning Groups (CPGs), and thin CPGs.
IBM SAN Volume Controller (SVC)	Using SSH communication	Managed disks, internal drives, virtual disks, disk groups, thin pools, array ports, adapters, host associations, and replications (local and remote).

About the discovery host

In Veritas InfoScale Operations Manager, to discover the array information on the enclosures using the Storage Insight Add-on, you need to designate a managed host as the discovery host. A discovery host is any Windows, Linux, or Solaris (SPARC) host (with a required version of `VRTSsfmh` package) added to the Veritas InfoScale Operations Manager domain. The discovery host must have vendor-specific array management tools installed on it.

See “About the network-attached storage discovery” on page 54.

See “About performance metering statistics for enclosure” on page 99.

See “About discovery methods and discovered information of storage enclosures” on page 52.

See “About Storage Insight Add-on” on page 51.

About the network-attached storage discovery

You can use Storage Insight Add-on to discover the network-attached storage (NAS) objects from the NAS enabled-enclosures in your data center.

When you configure an enclosure using Storage Insight Add-on, you can enable NAS discovery. After you configure the enclosure, the NAS objects are discovered, and displayed on Veritas InfoScale Operations Manager console. Veritas InfoScale Operations Manager supports NAS discovery for the following enclosures:

- NetApp
- EMC Celerra
- EMC VNX (File)

See “About Storage Insight Add-on” on page 51.

See “About discovery methods and discovered information of storage enclosures” on page 52.

See “About the discovery host” on page 53.

Adding HITACHI storage enclosures for deep discovery

The Storage Insight Add-on provides you detailed information on the storage enclosures in your data center. The Storage Insight Add-on makes use of HiCommand server that is connected to a host, which you can designate as the discovery host to get the information from HITACHI enclosures. This host need not be a part of the Management Server. The HITACHI agentlet that resides on the Management Server pushes the application programming interface (API) commands to the discovery host to get the information from the HiCommand server using XML. The discovery host runs the commands to collect the information on the enclosures. This information, which is brought to the Management Server, is parsed and put in the Veritas InfoScale Operations Manager database.

Before you add an enclosure for deep discovery, ensure that you meet the configuration prerequisites.

See “HITACHI enclosure configuration prerequisites” on page 111.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To add HITACHI storage enclosures for deep discovery

1. In the Home page on the Management Server console, click **Settings**.
2. Do one of the following:
   - Click **Device** and click **Add Enclosure**.
   - Click **Add Enclosure**.
3. In the **Add Enclosure** wizard panel, select the vendor and the enclosure model. Click **Next**.

   See "Add Enclosure panel options for vendor and product selection" on page 55.
4. In the **Add Enclosure** wizard panel to specify the details of the devices, enter the required information. Click **Next**.

   See "Add Enclosure panel options for enclosure selection" on page 56.
5. In the **Add Enclosure** wizard panel that lists the HITACHI enclosures that are connected to the HiCommand server, enter the required information to update discovery information. Click **Finish**.

   See "Add Enclosure panel options to update the enclosures discovery information" on page 66.

See "HITACHI storage enclosure commands" on page 135.

See "About Storage Insight Add-on" on page 51.

See "About discovery methods and discovered information of storage enclosures" on page 52.

See "About the discovery host" on page 53.

Add Enclosure panel options for vendor and product selection

Use this wizard panel to select the vendor, and the product for which you want to enable the deep discovery using Storage Insight Add-on.

Table 10-2 Add Enclosure panel for vendor and product selection

Field	Description
Enclosure Vendor	Select the enclosure vendor from the drop-down list.
Select product with appropriate discovery method	

Table 10-2 Add Enclosure panel for vendor and product selection *(continued)*

Field	Description
Product	Select the array model for which you want to enable the deep array discovery.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Additional Information	Provides the information about the discovered objects, prerequisites, CLI version, and other details about the enclosure discovery.

See "Adding HITACHI storage enclosures for deep discovery" on page 54.

See "Adding IBM XIV storage enclosures for deep discovery" on page 69.

See "Adding EMC Symmetrix storage enclosures for deep discovery " on page 68.

See "Adding NetApp storage enclosures for deep discovery" on page 70.

See "Adding EMC CLARiiON storage enclosures for deep discovery" on page 72.

See "Adding HP EVA storage enclosures for deep discovery" on page 73.

See "Adding EMC Celerra storage enclosures for deep discovery" on page 75.

See "Adding IBM System Storage DS enclosures for deep discovery" on page 74.

See "Adding EMC VNX storage enclosures for deep discovery" on page 76.

See "Adding EMC VPLEX storage enclosures for deep discovery" on page 78.

See "Adding 3PAR storage enclosures for deep discovery" on page 79.

See "Adding IBM SVC storage enclosures for deep discovery" on page 80.

Add Enclosure panel options for enclosure selection

Use this wizard panel to specify the details of the devices, and the server information for adding the following enclosures for deep discovery.

Table 10-3 lists the options for HITACHI enclosures.

Table 10-4 lists the options for EMC Symmetrix enclosures.

Table 10-5 lists the options for IBM XIV enclosures.

Table 10-6 lists the options for NetApp enclosures.

Table 10-7 lists the options for EMC CLARiiON enclosures.

Table 10-8 lists the options for HP Enterprise Virtual Array (EVA) enclosures.

Table 10-9 lists the options for IBM System Storage DS enclosures.

Table 10-10 lists the options for EMC Celerra enclosures.

Table 10-11 lists the options for EMC VNX enclosures.

Table 10-12 lists the options for EMC VPLEX enclosures.

Table 10-13 lists the options for 3PAR enclosures.

Table 10-14 lists the options for IBM SVC enclosures.

Note: Refer to the enclosure configuration prerequisites section to learn more about array configuration requirements.

Table 10-3 Add Enclosure panel options for HITACHI enclosure

Field	Description
Discovery Host	Enter the name of the managed host. A managed host that runs on Linux, Windows, or Solaris (SPARC) can be designated as a discovery host.
HiCommand Server	Enter the name of the HiCommand server.
Port	The port number of the HiCommand server. By default, the port number is 2001.
Username	Enter the user name for the HiCommand server.
Password	Enter the password for the HiCommand server.
Protocol	Enter either HTTP or HTTPS.

See "Adding HITACHI storage enclosures for deep discovery" on page 54.

Table 10-4 Add Enclosure panel options for EMC Symmetrix enclosure

Field	Description
Discovery Host	Enter the name of the managed host. A managed host that runs on Linux, Windows, or Solaris (SPARC) can be designated as a discovery host.

Table 10-4 Add Enclosure panel options for EMC Symmetrix enclosure *(continued)*

Field	Description
SYMAPI Server	Specify the SYMAPI Server name that is configured on the discovery host to discover the EMC Symmetrix enclosures. Use this option if discovery host does not have visibility to gatekeeper devices for Symmetrix enclosures.
SymCLI Location	
Use Default	Choose this option if you have stored the SymCLI binaries on the default location. SymCLI must be functional to discover the arrays so that Veritas InfoScale Operations Manager Management Server can provide details. Refer to the Enclosure configurations prerequisites section for more details.
Custom	Choose this option if you have stored the SymCLI binaries on any other location. Enter the path to the location in the field. You must ensure that the SymCLI binaries are available on the discovery host.
Enable performance metering	Enables the performance metering for the enclosure. Clear the check box to disable performance metering.

See “Adding EMC Symmetrix storage enclosures for deep discovery ” on page 68.

Table 10-5 Add Enclosure panel options for IBM XIV enclosures

Field	Description
Discovery Host	Enter the name of the managed host. A managed host that runs on Linux, Windows, or Solaris (SPARC) can be designated as a discovery host.
IBM XIV Server Name/IP 1	The IP address for the XIV system. The Storage Insight Add-on uses this IP address to connect to the XIV system using the XCLI utility.

Table 10-5 Add Enclosure panel options for IBM XIV enclosures *(continued)*

Field	Description
IBM XIV Server Name/IP 2	The IP address for the XIV system. The Storage Insight Add-on uses this IP address to connect to the XIV system using the XCLI utility, if the IP address that you have specified in the previous field is not accessible.
IBM XIV Server Name/IP 3	The IP address for the XIV system. The Storage Insight Add-on uses this IP address to connect to the XIV system using the XCLI utility, if the IP addresses that you have specified in the previous two fields are not accessible.
Username	Enter the user name for the XIV system.
Password	Enter the password for the XIV system.
XCLI Location	The location of the XCLI utility on the discovery host. You must ensure that the XCLI utility is available on the discovery host.
Enable performance metering	Enables the performance metering for the enclosure. Clear the check box to disable performance metering.

See “Adding IBM XIV storage enclosures for deep discovery” on page 69.

Table 10-6 Add Enclosure panel options for NetApp enclosures

Discovery Host	Enter the name of the managed host. A managed host that runs on Linux, Windows, or Solaris (SPARC) can be designated as a discovery host.
NetApp Server	Enter the name or the IP address for the NetApp server.

Table 10-6 Add Enclosure panel options for NetApp enclosures *(continued)*

Port	The port for the NetApp server. Enter 80 for communicating to the NetApp server over HTTP. For communicating to the NetApp server over HTTPS, enter 443. Ensure the port that you specify here is enabled on the enclosure.
Username	Enter the user name for the enclosure.
Password	Enter the password for the enclosure.
Enable NAS discovery	Select the check box to enable the NAS discovery for the NetApp enclosure.
Enable performance metering	Enables the performance metering for the enclosure. Clear the check box to disable performance metering.

See "Adding NetApp storage enclosures for deep discovery" on page 70.

Table 10-7 Add Enclosure panel options for EMC CLARiiON enclosures

Discovery Host	Enter the name of the managed host. A managed host that runs on Linux, Windows, or Solaris (SPARC) can be designated as a discovery host.
Storage Array Name/IP	Name or the IP address of the storage enclosure.
Port	The port for communicating to the EMC CLARiiON enclosure for getting the information. The default port number is 443.
Username	Enter the user name for the EMC CLARiiON enclosure.
Password	Enter the password for the EMC CLARiiON enclosure.

Table 10-7 Add Enclosure panel options for EMC CLARiiON enclosures *(continued)*

Scope	Specifies the type of the user account on the storage system that you want to log on. The available options are: ■ Global: Choose this option if your account is effective throughout the domain. When the administrator creates a global account, the software copies the definition of this account to the domain directory, which makes it accessible on all storage systems in the domain. ■ Local: Choose this option if your account is effective only on the storage systems for which the administrator creates the account. Using the local account, you can log on to only those storage systems on which you have a local account. ■ LDAP: LDAP maps the user name and the password entries to an external LDAP or Active Directory server for authentication. The user name and the password pairs whose roles are not mapped to the external directory are denied access. For authentication within the local security directory, specify global or local scope.
NAVISEC CLI Location	The location of the NaviSecCLI binary in the discovery host.
Use Default	Choose this option if you have stored the NaviSecCLI binaries on the default location. You must ensure that the NaviSecCLI binaries are available on the discovery host.
Custom	Choose this option if you have stored the NaviSecCLI binaries on any other location. Enter the path to the location in the field. You must ensure that the NaviSecCLI binaries are available on the discovery host.

Table 10-7 Add Enclosure panel options for EMC CLARiiON enclosures *(continued)*

Use Secure Socket Layer	Select this check box to use the secure socket layer for communicating to the enclosure. If you select this check box, you do not have to enter the credentials again when you perform deep discovery for the EMC CLARiiON enclosures.
Enable performance metering	Enables the performance metering for the enclosure. Clear the check box to disable performance metering.

See "Adding EMC CLARiiON storage enclosures for deep discovery" on page 72.

Table 10-8 Add Enclosure panel options for HP EVA enclosure

Field	Description
Discovery Host	Enter the name of the managed host. A managed host that runs on Linux, Windows, or Solaris (SPARC) can be designated as a discovery host.
Command View Server/IP	Enter the name, or the IP address of the Command View server.
User Name	Enter the user name for the Command View server. Note that along with the root or the admin user, a non-admin user can also perform the discovery.
Password	Enter the password for the Command View server. If you have selected the password file option, you do not need to enter the password.
SSSU CLI Location	Enter the full path for the Storage Scripting System Utility (SSSU) CLI on the discovery host.

Table 10-8 Add Enclosure panel options for HP EVA enclosure *(continued)*

Field	Description
SSSU Password-File Directory Path	Enter the directory path for the SSSU password file on the discovery host. This file is generated using the SSSU utility. If you have already specified the password in the **Password** field, this entry is ignored.

See "Adding HP EVA storage enclosures for deep discovery" on page 73.

Table 10-9 Add Enclosure panel options for IBM System Storage DS enclosure

Field	Description
Discovery Host	Enter the name of the managed host. A managed host that runs on Linux, Windows, or Solaris (SPARC) can be designated as a discovery host. This host must have DSCLI utility installed on it. It should also be connected to the enclosure.
Storage array name/IP	Enter the name, or the IP address for the IBM System Storage DS enclosure.
User Name	Enter the user name for the enclosure.
Password	Enter the password to access the enclosure. Leave this field blank if the discovery is done using the password file.
DSCLI Location	Enter the full directory path of the DSCLI utility on the discovery host.
Password File Path	Enter the full path (including the file name) for the password file on the discovery host. This field is ignored if the password is provided in the **Password** field.

See "Adding IBM System Storage DS enclosures for deep discovery" on page 74.

Table 10-10 Add Enclosure panel options for EMC Celerra enclosure

Field	Description
Discovery Host	Enter the name of the managed host. A managed host that runs on Linux, Windows, or Solaris (SPARC) can be designated as a discovery host.
Celerra Control Station	Interface for Celerra communication.
Username	Enter the user name for the enclosure.
Password	Enter the password for the enclosure.

See "Adding EMC Celerra storage enclosures for deep discovery" on page 75.

Table 10-11 Add Enclosure panel options for EMC VNX enclosure

Field	Description
Discovery Host	Enter the name of the managed host. A managed host that runs on Linux, Windows, or Solaris (SPARC) can be designated as a discovery host.
Control Station IP/Name	Enter the name or the IP address of the VNX File. It is optional for VNX Block only device.
Block IP/Name:Scope	IP address or name and scope for VNX Block. The IP address or name, and the scope must be separated using colon. It is optional for VNX File only device.
User Name	Enter the user name for the enclosure.
Password	Enter the password for the enclosure.
CLI Location	Enter the Navisphere CLI location.
Enable performance metering	Enables the performance metering for the enclosure. Clear the check box to disable performance metering.

See "Adding EMC VNX storage enclosures for deep discovery" on page 76.

Table 10-12 Add Enclosure panel options for EMC VPLEX enclosure

Field	Description
Discovery Host	Enter the name of the managed host. A managed host that runs on Linux, Windows, or Solaris (SPARC) can be designated as a discovery host.
VPLEX Array Name/IP	Enter the name, or the IP address of the VPLEX enclosure.
User Name	Enter the user name for the enclosure. To perform the enclosure discovery, you must have administrator privileges.
Password	Enter the password for the enclosure.
Port	Enter the communication port for the VPLEX enclosure. The default communication port is 443.

See "Adding EMC VPLEX storage enclosures for deep discovery" on page 78.

Table 10-13 Add Enclosure panel options for 3PAR enclosure

Field	Description
Discovery Host	Enter the name of the managed host. A managed host that runs on Linux, Windows, or Solaris (SPARC) can be designated as a discovery host.
3PAR Enclosure IP/Name	Enter the name or the IP address for the 3PAR enclosure.
User Name	Enter the user name for the enclosure.
Password	Enter the password for the enclosure.

See "Adding 3PAR storage enclosures for deep discovery" on page 79.

Table 10-14 Add Enclosure panel options for IBM SVC enclosure

Field	Description
Discovery Host	Enter the name of the managed host. A managed host that runs on Linux, Windows, or Solaris (SPARC) can be designated as a discovery host.
Enclosure IP/Name	Enter the name, or the IP address of the IBM SVC enclosure.
User Name	Enter the user name for the enclosure. Specify admin as the username for IBM SVC enclosure with version 5.1.
Password	Enter the password for the enclosure. Leave this field blank if the discovery is done using the ssh private key file.
Certificate Location	Enter the full path (including the file name) for the ssh private key file on the discovery host.

See "Adding IBM SVC storage enclosures for deep discovery" on page 80.

Add Enclosure panel options to update the enclosures discovery information

Use this wizard panel to update the deep discovery information for an enclosure.

Table 10-15 Add Enclosure panel to update the deep discovery information for enclosures

Field	Description
Configuration Name	Enter a name for the deep discovery operation that you want to perform.
Enclosures	
Display Name	Displays the name of the enclosure.
Vendor ID	Displays the ID that Veritas InfoScale Operations Manager generates for the enclosure.
Serial	Displays the serial number of the enclosure.

Table 10-15 Add Enclosure panel to update the deep discovery information for enclosures *(continued)*

Field	Description
Vendor	Displays manufacturer of the enclosure.
Model	Displays the enclosure model information.
Product	Displays the type of the enclosure.
Connectivity	This field is displayed only for the EMC Symmetrix enclosures. It indicates the following: ■ **Local**: Whether the enclosure is connected to the host locally. ■ **Remote**: Whether the enclosure is connected to another enclosure, using the EMC Symmetrix Remote Data Facility (SRDF), which might have been connected to the host locally.
Discovery	Choose **Enable** to perform the deep discovery. Choose **Disable** to disable the deep discovery.
Configured Name	If the enclosure was already add for deep discovery, the configuration name that was entered at that time is displayed in this field.

Select the check box in the top row to enable deep discovery for all the enclosures in the list.

See “Adding HITACHI storage enclosures for deep discovery” on page 54.

See “Adding IBM XIV storage enclosures for deep discovery” on page 69.

See “Adding EMC Symmetrix storage enclosures for deep discovery ” on page 68.

See “Adding NetApp storage enclosures for deep discovery” on page 70.

See “Adding EMC CLARiiON storage enclosures for deep discovery” on page 72.

See “Adding HP EVA storage enclosures for deep discovery” on page 73.

See “Adding EMC Celerra storage enclosures for deep discovery” on page 75.

See “Adding IBM System Storage DS enclosures for deep discovery” on page 74.

See “Adding EMC VNX storage enclosures for deep discovery” on page 76.

See "Adding EMC VPLEX storage enclosures for deep discovery" on page 78.

See "Adding 3PAR storage enclosures for deep discovery" on page 79.

See "Adding IBM SVC storage enclosures for deep discovery" on page 80.

Adding EMC Symmetrix storage enclosures for deep discovery

The Storage Insight Add-on provides you detailed information on the storage enclosures in your data center. It uses SymCLI utilities that communicate to the Symmetrix enclosures in your data center. The SymCLI utilities must be configured on the discovery host, which must be part of the Veritas InfoScale Operations Manager domain.

You can perform deep discovery also for the Symmetrix enclosures that are connected the discovery hosts using the EMC Symmetrix Remote Data Facility (SRDF).

After you add EMC Symmetrix enclosures for deep discovery, the discovery happens once every six hours. After the initial discovery, only add, remove, commit, resume, suspend, or split actions on the following are captured in the subsequent discoveries:

- Snapshots
- Configuration changes
- Device masking
- BCV
- Clones
- RDF and FAST VP

You must use SymCLI Version 7.0 or later to perform deep discovery on the EMC Symmetrix storage enclosures. To discover the Media Type and the RPM of disks in a Symmetrix enclosure, you must use SymCLI Version 7.1 or later.

For more information on SymCLI versions for discovering FAST managed storage groups in a Symmetrix VMAX enclosure, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To add EMC Symmetrix storage enclosures for deep discovery

1. In the Home page on the Management Server console, click **Settings**.
2. Do one of the following:

- Click **Device** and click **Add Enclosure**.
- Click **Add Enclosure**.

3 In the **Add Enclosure** wizard panel, select the vendor and the enclosure model. Click **Next.**

See "Add Enclosure panel options for vendor and product selection" on page 55.

4 In the **Add Enclosure** wizard panel to specify the details the deep discovery, enter the required information. Click **Next**.

See "Add Enclosure panel options for enclosure selection" on page 56.

5 In the **Add Enclosure** wizard panel to update discovery information, choose the enclosures and enable the deep array discovery. Click **Finish**.

See "Add Enclosure panel options to update the enclosures discovery information" on page 66.

See "EMC Symmetrix storage enclosure commands" on page 135.

See "About Storage Insight Add-on" on page 51.

See "About discovery methods and discovered information of storage enclosures" on page 52.

See "About performance metering statistics for enclosure" on page 99.

Adding IBM XIV storage enclosures for deep discovery

The Storage Insight Add-on provides you detailed information on the storage enclosures in your data center. It makes use of the XCLI utility that communicates to the IBM XIV systems using the TCP/IP protocol. The XCLI utility must be configured on the discovery host, which must be part of the Veritas InfoScale Operations Manager domain.

Note: The size of 1 MB in an IBM XIV enclosure is equivalent to 1000 KB, unlike the normal 1024 KB. Therefore, the disk capacity may vary from what you actually see on an IBM XIV enclosure on the Veritas InfoScale Operations Manager console.

Before you add the IBM XIV enclosure for deep discovery of enclosure, ensure that you meet the configuration prerequisites.

See "IBM XIV enclosure configuration prerequisites" on page 117.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To add the IBM XIV storage enclosures for deep discovery

1. In the Home page on the Management Server console, click **Settings**.
2. Do one of the following:
   - Click **Device** and click **Add Enclosure**.
   - Click **Add Enclosure**.
3. In the **Add Enclosure** wizard panel, select the vendor and the enclosure model. Click **Next.**

   See "Add Enclosure panel options for vendor and product selection" on page 55.
4. In the **Add Enclosure** wizard panel to specify the details of the devices, enter the required information. Click **Next**.

   See "Add Enclosure panel options for enclosure selection" on page 56.
5. In the **Add Enclosure** wizard panel that lists the IBM XIV enclosures that are connected to the XCLI utility, enter the required information to update discovery information. Click **Finish**.

   See "Add Enclosure panel options to update the enclosures discovery information" on page 66.

See "IBM XIV storage enclosures commands" on page 137.

See "About Storage Insight Add-on" on page 51.

See "About discovery methods and discovered information of storage enclosures" on page 52.

See "About performance metering statistics for enclosure" on page 99.

Adding NetApp storage enclosures for deep discovery

The Storage Insight Add-on provides you detailed information on the storage enclosures in your data center. The NetApp storage objects work on the Data ONTAP operating system, which provides various interfaces to administer the NetApp storage objects. Veritas InfoScale Operations Manager communicates to the enclosures using the ONTAP SDK interface to get the NetApp enclosure information. This communication occurs through the HTTP protocol (using the port number 80) or through the HTTPS protocol (using the port number 443).

It is recommended that you must configure all the nodes for deep discovery in a NetApp cluster.

After you add the NetApp enclosure for deep discovery, you can view the following information on the details page:

- Details of the enclosure such as the model, serial number, version of the Data ONTAP installed on the enclosure, the physical capacity of the enclosure, number of spare disks and the name of the partner node
- The physical disks on the enclosure
- Aggregates
- Flex/traditional volumes
- LUNs
- Array ports
- Replications
- NAS objects: Qtree, Share, Quota, vFiler, and NAS/CIFS share consumer information

---

Note: To create both the flexible and the traditional volumes on the aggregates, you must use Data ONTAP version 7.0 or later. Traditional volumes can exist on the enclosures that have Data ONTAP 6.0 or later. Storage Insight Add-on supports the NetApp enclosures that have Data ONTAP 6.5.2 or later.

Before you add the NetApp enclosure for deep discovery, ensure that you meet the configuration prerequisites.

See "NetApp storage enclosure configuration prerequisites" on page 118.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To add the NetApp enclosures for deep discovery

1. In the Home page on the Management Server console, click **Settings**.
2. Do one of the following:
   - Click **Device** and click **Add Enclosure**.
   - Click **Add Enclosure**.
3. In the **Add Enclosure** wizard panel, select the vendor and the enclosure model. Click **Next.**

   See "Add Enclosure panel options for vendor and product selection" on page 55.

4 In the **Add Enclosure** wizard panel to specify the details of the devices, enter the required information. Click **Next**.

See "Add Enclosure panel options for enclosure selection" on page 56.

5 In the **Add Enclosure** wizard panel that lists the NetApp enclosures, enter the required information to update discovery information. Click **Finish**.

See "Add Enclosure panel options to update the enclosures discovery information" on page 66.

See "NetApp storage enclosure commands" on page 139.

See "About Storage Insight Add-on" on page 51.

See "About discovery methods and discovered information of storage enclosures" on page 52.

See "About performance metering statistics for enclosure" on page 99.

See "About the discovery host" on page 53.

Adding EMC CLARiiON storage enclosures for deep discovery

The Storage Insight Add-on provides you detailed information on the storage enclosures in your data center. Storage Insight Add-on communicates to the EMC CLARiiON enclosures through the NaviSphere Secure CLI (NaviSecCLI) utility, which is the secure CLI for communicating to the NaviSphere package on the enclosure. The NaviSphere Secure CLI utility must be installed on a discovery host that is a part of the Veritas InfoScale Operations Manager domain.

Storage Insight Add-on supports only the EMC CLARiiON CX series enclosures with a firmware version 6.26 and later.

Before you add the enclosure for deep discovery, ensure that you meet the configuration prerequisites.

See "EMC CLARiiON storage enclosures configuration prerequisites" on page 119.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To add the EMC CLARiiON enclosures for deep discovery

1 In the Home page on the Management Server console, click **Settings**.

2 Do one of the following:

- Click **Device** and click **Add Enclosure**.
- Click **Add Enclosure**.

3 In the **Add Enclosure** wizard panel, select the vendor and the enclosure model. Click **Next**.

See "Add Enclosure panel options for vendor and product selection" on page 55.

4 In the **Add Enclosure** wizard panel to specify the details of the devices, enter the required information. Click **Next**.

See "Add Enclosure panel options for enclosure selection" on page 56.

5 In the **Add Device** wizard panel that lists the EMC CLARiiON enclosures, enter the required information to update discovery information. Click **Finish**.

See "Add Enclosure panel options to update the enclosures discovery information" on page 66.

See "EMC CLARiiON storage enclosure commands" on page 140.

See "About Storage Insight Add-on" on page 51.

See "About discovery methods and discovered information of storage enclosures" on page 52.

See "About performance metering statistics for enclosure" on page 99.

See "About the discovery host" on page 53.

Adding HP EVA storage enclosures for deep discovery

The Storage Insight Add-on provides you the detailed information about the configured HP Enterprise Virtual Arrays (EVA) in your data center. It includes the information about virtual disks, physical disks, array disk groups, and replications. Veritas InfoScale Operations Manager Management Server communicates with the HP Command View EVA software that manages the HP EVA storage arrays.

Before you add the enclosure for deep discovery, ensure that you meet the configuration prerequisites.

See "Hewlett-Packard Enterprise Virtual Array (HP EVA) configuration prerequisites" on page 122.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To add the HP EVA storage enclosures for deep discovery

1 In the Home page on the Management Server console, click **Settings**.

2 Do one of the following:

- Click **Device** and click **Add Enclosure**.

- Click **Add Enclosure**.

3 In the **Add Enclosure** wizard panel, select the vendor, and the enclosure model. Click **Next.**

See "Add Enclosure panel options for vendor and product selection" on page 55.

4 In the **Add Enclosure** wizard panel, enter the details for the HP EVA device. Click **Next**.

See "Add Enclosure panel options for enclosure selection" on page 56.

5 In the **Add Enclosure** wizard panel that lists the HP EVA enclosures that are connected to the HP Command View utility, enter the required information to update the enclosure discovery. Click **Finish**.

See "Add Enclosure panel options to update the enclosures discovery information" on page 66.

See "HP EVA storage enclosure commands" on page 141.

See "About Storage Insight Add-on" on page 51.

See "About discovery methods and discovered information of storage enclosures" on page 52.

See "About the discovery host" on page 53.

Adding IBM System Storage DS enclosures for deep discovery

You can use Storage Insight Add-on to perform the deep discovery of IBM System Storage DS enclosures.

For IBM System Storage DS enclosure, you might observe inconsistency between the vendor data and Veritas InfoScale Operations Manager storage calculation. This is because Veritas InfoScale Operations Manager displays all sizes in the multiples of 1024. For example:

- 1 kilobyte (KB) = 1,024 bytes
- 1 megabyte (MB) = 2^20 bytes = 1,048,576 bytes
- 1 gigabyte (GB) = 2^30 bytes = 1,024 MB = 1,073,741,824 bytes
- 1 terabyte (TB) = 2^40 bytes = 1,024 GB = 1,099,511,627,776 bytes

The storage vendor displays sizes in the multiples of 1024 bytes, or 1000 bytes.

Before you add the enclosure for deep discovery, ensure that you meet the configuration prerequisites.

See "IBM System Storage DS enclosure configuration prerequisites" on page 125.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To add the IBM System Storage DS enclosures for deep discovery

1. In the Home page on the Management Server console, click **Settings**.
2. Do one of the following:
   - Click **Device** and click **Add Enclosure**.
   - Click **Add Enclosure**.
3. In the **Add Enclosure** wizard panel, select the vendor and the enclosure model. Click **Next**.

   See "Add Enclosure panel options for vendor and product selection" on page 55.
4. In the **Add Enclosure** wizard panel, enter the required information for the IBM System Storage DS enclosure. Click **Next**.

   See "Add Enclosure panel options for enclosure selection" on page 56.
5. In the **Add Enclosure** wizard panel that lists the IBM System Storage DS enclosures, enter the required information to update the enclosure discovery. Click **Finish**.

   See "Add Enclosure panel options to update the enclosures discovery information" on page 66.

See "IBM System Storage DS enclosure commands" on page 142.

See "About Storage Insight Add-on" on page 51.

See "About discovery methods and discovered information of storage enclosures" on page 52.

See "About the discovery host" on page 53.

Adding EMC Celerra storage enclosures for deep discovery

The Storage Insight Add-on provides you detailed information on the storage enclosures in your data center. To configure deep discovery, you must have at least read-only permissions to query the Celerra array.

Before you add the enclosure for deep discovery, ensure that you meet the configuration prerequisites.

See "EMC Celerra enclosure configuration prerequisites" on page 127.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To add EMC Celerra storage enclosures for deep discovery

1. In the Home page on the Management Server console, click **Settings**.
2. Do one of the following:
   - Click **Device** and click **Add Enclosure**.
   - Click **Add Enclosure**.
3. In the **Add Enclosure** wizard panel, select the vendor, and the enclosure model. Click **Next.**

   See "Add Enclosure panel options for vendor and product selection" on page 55.
4. In the **Add Enclosure** wizard panel, enter the details for the EMC Celerra device. Click **Next**.

   See "Add Enclosure panel options for enclosure selection" on page 56.
5. In the **Add Enclosure** wizard panel that lists the EMC Celerra enclosures, enter the discovery information. Click **Finish**.

   See "Add Enclosure panel options to update the enclosures discovery information" on page 66.

See "EMC Celerra storage enclosure commands" on page 145.

See "About Storage Insight Add-on" on page 51.

See "About discovery methods and discovered information of storage enclosures" on page 52.

See "About the discovery host" on page 53.

Adding EMC VNX storage enclosures for deep discovery

You can use Storage Insight Add-on to perform the deep discovery of EMC VNX enclosures. After the Storage Insight Add-on performs the deep array discovery, you can view the following additional information on the overview page of the EMC VNX enclosure:

For EMC VNX (File):

- LUNs
- Disk Volumes
- Storage Pools

- Array Volumes
- Replications
- NAS file systems
- Shares
- Data Movers

For EMC VNX (Block):

- LUNs
- Physical Disks
- RAID Groups
- Replications
- Thin Pools

Before you add the enclosure for deep discovery, ensure that you meet the configuration prerequisites.

See "EMC VNX storage enclosure configuration prerequisites" on page 131.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To add the EMC VNX enclosures for deep discovery

1. In the Home page on the Management Server console, click **Settings**.
2. Do one of the following:
   - Click **Device** and click **Add Enclosure**.
   - Click **Add Enclosure**.
3. In the **Add Enclosure** wizard panel, select the vendor and the enclosure model. Click **Next**.

   See "Add Enclosure panel options for vendor and product selection" on page 55.
4. In the **Add Enclosure** wizard panel, enter the required information about the EMC VNX enclosure. Click **Next**.

   See "Add Enclosure panel options for enclosure selection" on page 56.
5. In the **Add Enclosure** wizard panel that lists the EMC VNX enclosures, enter the deep array discovery information. Click **Finish**.

   See "Add Enclosure panel options to update the enclosures discovery information" on page 66.

See "EMC VNX (Block) storage enclosure commands" on page 146.

See "EMC VNX (File) storage enclosure commands" on page 147.

See "About Storage Insight Add-on" on page 51.

See "About discovery methods and discovered information of storage enclosures" on page 52.

See "About performance metering statistics for enclosure" on page 99.

See "About the discovery host" on page 53.

Adding EMC VPLEX storage enclosures for deep discovery

You can use Storage Insight Add-on to perform the deep discovery of EMC VPLEX enclosures.

A VPLEX enclosure is represented as a virtual container object that contains one or more VPLEX clusters. The container object's waterfall page provides the information about the objects seen by all nodes. The individual tabs display the details of the distributed objects. The nodes show the details of the objects that can be seen only from the nodes.

Note: To add a VPLEX enclosure, the user account used for the discovery should have the administrator privileges on the VPLEX enclosure.

Before you add the enclosure for deep discovery, ensure that you meet the configuration prerequisites.

See "EMC VPLEX storage enclosure configuration prerequisites" on page 132.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To add EMC VPLEX enclosures for deep discovery

1. In the Home page on the Management Server console, click **Settings**.
2. Do one of the following:
   - Click **Device** and click **Add Enclosure**.
   - Click **Add Enclosure**.
3. In the **Add Enclosure** wizard panel, select the vendor and the enclosure model. Click **Next**.

   See "Add Enclosure panel options for vendor and product selection" on page 55.

4 In the **Add Enclosure** wizard panel, enter the required information about the EMC VPLEX enclosure. Click **Next**.

See "Add Enclosure panel options for enclosure selection" on page 56.

5 In the **Add Enclosure** wizard panel that lists the EMC VPLEX enclosures, enter the deep array discovery information. Click **Finish**.

See "Add Enclosure panel options to update the enclosures discovery information" on page 66.

See "EMC VPLEX storage enclosure commands" on page 148.

See "About Storage Insight Add-on" on page 51.

See "About discovery methods and discovered information of storage enclosures" on page 52.

See "About the discovery host" on page 53.

Adding 3PAR storage enclosures for deep discovery

You can use Storage Insight Add-on to perform the deep discovery of 3PAR enclosures.

If 3PAR enclosure is used with VMware ESX server, correct storage correlation might not be shown on Veritas InfoScale Operations Manager console. To discover correct storage correlation, you need to change the ESX server's persona from Generic to Generic-Legacy on 3PAR array.

Before you add the enclosure for deep discovery, ensure that you meet the configuration prerequisites.

See "3PAR storage enclosure configuration prerequisites" on page 133.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To add the 3PAR enclosures for deep discovery

1 In the Home page on the Management Server console, click **Settings**.

2 Do one of the following:

- Click **Device** and click **Add Enclosure**.
- Click **Add Enclosure**.

3 In the **Add Enclosure** wizard panel, select the vendor, and the enclosure model. Click **Next.**

See "Add Enclosure panel options for vendor and product selection" on page 55.

4 In the **Add Enclosure** wizard panel, enter the required information about the 3PAR enclosure. Click **Next**.

See "Add Enclosure panel options for enclosure selection" on page 56.

5 In the **Add Enclosure** wizard panel that lists the 3PAR enclosures, enter the deep array discovery information. Click **Finish**.

See "Add Enclosure panel options to update the enclosures discovery information" on page 66.

See "3PAR storage enclosure commands" on page 149.

See "About Storage Insight Add-on" on page 51.

See "About discovery methods and discovered information of storage enclosures" on page 52.

See "About the discovery host" on page 53.

Adding IBM SVC storage enclosures for deep discovery

You can use Storage Insight Add-on to perform the deep discovery of IBM SVC enclosures.

Before you add the enclosure for deep discovery, ensure that you meet the configuration prerequisites.

See "IBM SVC enclosure configuration prerequisites" on page 126.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To add the IBM SVC enclosures for deep discovery

1 In the Home page on the Management Server console, click **Settings**.

2 Do one of the following:

- Click **Device** and click **Add Enclosure**.
- Click **Add Enclosure**.

3 In the **Add Enclosure** wizard panel, select the vendor and the enclosure model. Click **Next**.

See "Add Enclosure panel options for vendor and product selection" on page 55.

4 In the **Add Enclosure** wizard panel, enter the required information for the IBM SVC enclosure. Click **Next**.

See “Add Enclosure panel options for enclosure selection” on page 56.

5 In the **Add Enclosure** wizard panel that lists the enclosures, enter the required information to update the enclosure discovery. Click **Finish**.

See “Add Enclosure panel options to update the enclosures discovery information” on page 66.

See “IBM SVC storage enclosure commands” on page 149.

See “About Storage Insight Add-on” on page 51.

See “About discovery methods and discovered information of storage enclosures” on page 52.

See “About the discovery host” on page 53.

Editing the deep discovery configuration for an enclosure

Using Storage Insight Add-on, you can modify the deep array discovery configuration information for an enclosure.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To edit the deep array discovery configuration for an enclosure

1 In the Home page on the Management Server console, click **Settings**.

2 Click **Device**.

3 Expand **Enclosures** to locate the vendor.

4 In the vendor configurations list, right-click the enclosure and select **Edit Configuration**.

5 In the **Edit Configuration** wizard panel, edit the configuration details to change the device discovery. Click **Next**.

See “Edit Configuration panel options to modify the device discovery” on page 82.

6 In the **Edit Configuration** wizard panel, select the check box for the enclosure for which you want to perform the deep discovery configuration. Click **Finish**.

See "Edit Configuration panel options to select an enclosure for deep discovery" on page 93.

7 In the **Edit Configuration** result panel review the information and click **OK**.

See "About Storage Insight Add-on" on page 51.

See "About the discovery host" on page 53.

Edit Configuration panel options to modify the device discovery

Use this wizard panel to edit the configuration for deep discovery for the following enclosures:

Table 10-16 lists the options for HITACHI enclosures.

Table 10-17 lists the options for EMC Symmetrix enclosures.

Table 10-18 lists the options for EMC Symmetrix enclosures.

Table 10-19 lists the options for NetApp enclosures.

Table 10-20 lists the options for EMC CLARiiON enclosures.

Table 10-21 lists the options for HP Enterprise Virtual Array (EVA) enclosures.

Table 10-22 lists the options for IBM System Storage DS enclosures.

Table 10-23 lists the options for EMC Celerra enclosures.

Table 10-24 lists the options for EMC VNX enclosures.

Table 10-25 lists the options for EMC VPLEX enclosures.

Table 10-26 lists the options for 3PAR enclosures.

Table 10-27 lists the options for IBM SVC enclosures.

Table 10-16 Edit configuration panel options for HITACHI enclosure

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.

Table 10-16 Edit configuration panel options for HITACHI enclosure *(continued)*

Field	Description
Discovery Host	Modify the name of the host. This host must be a part of the Veritas InfoScale Operations Manager Management Server domain.
HiCommand Server	Modify the name of the HiCommand server.
Port	Modify the port number of the HiCommand server. By default, the port number is 2001.
Username	Modify the user name if you have modified the name of the HiCommand server.
Password	Modify the password if you have modified the name of the HiCommand server.

Table 10-17 Edit configuration panel options for EMC Symmetrix enclosure

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the host. This host must be a part of the Veritas InfoScale Operations Manager Management Server domain.
SYMAPI Server	Specify the SYMAPI Server name that is configured on the discovery host to discovery EMC Symmetrix enclosures. Use this option if discovery host does not have visibility to gatekeeper devices for Symmetrix enclosures.
SymCLI Location	
Use Default	Choose this option if you have stored the SymCLI binaries on the default location. You must ensure that the SymCLI binaries are available on the discovery host.

Table 10-17 Edit configuration panel options for EMC Symmetrix enclosure *(continued)*

Field	Description
Custom	Choose this option if you have stored the SymCLI binaries on any other location. Enter the path to the location in the field. You must ensure that the SymCLI binaries are available on the discovery host.
Enable performance metering	Enables the performance metering for the enclosure. Clear the check box to disable performance metering.

Table 10-18 Edit configuration panel options for IBM XIV enclosures

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the host. This host must be a part of the Veritas InfoScale Operations Manager Management Server domain.
IBM XIV Server Name/IP 1	The IP address for the XIV system. The Storage Insight Add-on uses this IP address to connect to the XIV system using the XCLI utility.
IBM XIV Server Name/IP 2	The IP address for the XIV system. The Storage Insight Add-on uses this IP address to connect to the XIV system using the XCLI utility, if the IP address that you have specified in the previous field is not accessible

Table 10-18 Edit configuration panel options for IBM XIV enclosures *(continued)*

Field	Description
IBM XIV Server Name/IP 3	The IP address for the XIV system. The Storage Insight Add-on uses this IP address to connect to the XIV system using the XCLI utility, if the IP addresses that you have specified in the previous two fields are not accessible
Username	Modify the user name for the XIV system.
Password	Modify the password for the XIV system.
XCLI Location	The location of the XCLI utility on the discovery host. You must ensure that the XCLI utility is available on the discovery host.
Enable performance metering	Enables the performance metering for the enclosure. Clear the check box to disable performance metering.

Table 10-19 Edit configuration panel options for NetApp enclosures

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the host. This host must be a part of the Veritas InfoScale Operations Manager Management Server domain.
NetApp Server	Modify the name or the IP address for the NetApp server.

Table 10-19 Edit configuration panel options for NetApp enclosures *(continued)*

Field	Description
Port	Modify the port for the NetApp server. Enter 80 for communicating to the NetApp server over HTTP. For communicating to the NetApp server over HTTPS, enter 443. Ensure the port that you specify here is enabled on the enclosure.
Username	Modify the user name for the enclosure.
Password	Modify the password for the enclosure.
Enable NAS discovery	Select the check box to enable the NAS discovery for the NetApp enclosure.
Enable performance metering	Enables the performance metering for the enclosure. Clear the check box to disable performance metering.

Table 10-20 Edit configuration panel options for EMC CLARiiON enclosures

Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the host. This host must be a part of the Veritas InfoScale Operations Manager Management Server domain.
Storage Array Name/IP	Modify the name or the IP address of the storage enclosure.
Port	The port for communicating to the EMC CLARiiON enclosure for getting the information. The default port number is 443.

Table 10-20 Edit configuration panel options for EMC CLARiiON enclosures *(continued)*

Username	Modify the user name for the EMC CLARiiON enclosure.
Password	Modify the password for the EMC CLARiiON enclosure.
Scope	Specifies the type of the user account on the storage system that you want to log on. The available options are: ■ Global: Choose this option if your account is effective throughout the domain. When the administrator creates a global account, the software copies the definition of this account to the domain directory, which makes it accessible on all storage systems in the domain. ■ Local: Choose this option if your account is effective only on the storage systems for which the administrator creates the account. Using the local account, you can log on to only those storage systems on which you have a local account. ■ LDAP: LDAP maps the user name and the password entries to an external LDAP or Active Directory server for authentication. The user name and the password pairs whose roles are not mapped to the external directory are denied access. For authentication within the local security directory, specify global or local scope.
NAVISEC CLI Location	The location of the NaviSecCLI binary in the discovery host.
Use Default	Choose this option if you have stored the NaviSecCLI binaries on the default location. You must ensure that the NaviSecCLI binaries are available on the discovery host.

Table 10-20 Edit configuration panel options for EMC CLARiiON enclosures *(continued)*

Custom	Choose this option if you have stored the NaviSecCLI binaries on any other location. Enter the path to the location in the field. You must ensure that the NaviSecCLI binaries are available on the discovery host.
Use Secure Socket Layer	Select this check box to use the secure socket layer for communicating to the enclosure. If you select this check box, you do not have to enter the credentials again when you perform deep discovery for the EMC CLARiiON enclosures.
Enable performance metering	Enables the performance metering for the enclosure. Clear the check box to disable performance metering.

Table 10-21 Edit configuration panel options for HP EVA enclosure

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the host. This host must be a part of the Veritas InfoScale Operations Manager Management Server domain.
Command View Server/IP	Modify the name, or the IP address of the Command View server.
User Name	Modify the user name for the Command View server if you have modified the name. Note that along with the root or the admin user, a non-admin user can also perform the discovery.

Table 10-21 Edit configuration panel options for HP EVA enclosure *(continued)*

Field	Description
Password	Modify the password for the Command View server. If you have selected the password file option, you do not need to enter the password.
SSSU CLI Location	Modify the full path for the Storage Scripting System Utility (SSSU) CLI on the discovery host.
SSSU Password-File Directory Path	Enter the directory path for the SSSU password file on the discovery host. This file is generated using the SSSU utility. If you have already specified the password in the **Password** field, this entry is ignored.

Table 10-22 Edit configuration panel options for IBM System Storage DS enclosure

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the host. This host must be a part of the Veritas InfoScale Operations Manager Management Server domain. This host must have DSCLI utility installed on it. It should also be connected to the enclosure.
Storage array name/IP	Modify the name, or the IP address for the IBM System Storage DS enclosure.
User Name	Modify the user name for the enclosure.
Password	Modify the password to access the enclosure. Leave this field blank if the discovery is done using the password file.
DSCLI Location	Modify the full directory path of the DSCLI utility on the discovery host.

Table 10-22 Edit configuration panel options for IBM System Storage DS enclosure *(continued)*

Field	Description
Password File Path	Enter the full path (including the file name) for the password file on the discovery host. This field is ignored if the password is provided in the **Password** field.

Table 10-23 Edit configuration panel options for EMC Celerra enclosure

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the host. This host must be a part of the Veritas InfoScale Operations Manager Management Server domain.
Celerra Control Station	Modify the interface for Celerra communication.
Username	Modify the user name for the enclosure.
Password	Modify the password for the enclosure.

Table 10-24 Edit configuration panel options for EMC VNX enclosure

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the host. This host must be a part of the Veritas InfoScale Operations Manager Management Server domain.

Table 10-24 Edit configuration panel options for EMC VNX enclosure *(continued)*

Field	Description
Control Station IP/Name	Modify the name or the IP address of the VNX File. It is optional for VNX Block only device.
Block IP/Name:Scope	Modify the IP address or name and scope for VNX Block. The IP address or name, and the scope must be separated using colon. It is optional for VNX File only device.
User Name	Modify the user name for the enclosure.
Password	Modify the password for the enclosure.
CLI Location	Modify the Navisphere CLI location.
Enable performance metering	Enables the performance metering for the enclosure. Clear the check box to disable performance metering.

Table 10-25 Edit configuration panel options for EMC VPLEX enclosure

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the host. This host must be a part of the Veritas InfoScale Operations Manager Management Server domain.
VPLEX Array Name/IP	Modify the name, or the IP address of the VPLEX enclosure.
User Name	Modify the user name for the enclosure. To perform the enclosure discovery, you must have administrator privileges.
Password	Modify the password for the enclosure.

Table 10-25 Edit configuration panel options for EMC VPLEX enclosure *(continued)*

Field	Description
Port	Modify the communication port for the VPLEX enclosure. The default communication port is 443.

Table 10-26 Edit configuration panel options for 3PAR enclosure

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the host. This host must be a part of the Veritas InfoScale Operations Manager Management Server domain.
3PAR Enclosure IP/Name	Modify the name or the IP address for the 3PAR enclosure.
User Name	Modify the user name for the enclosure.
Password	Modify the password for the enclosure.

Table 10-27 Edit configuration panel options for IBM SVC enclosure

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the discovery host. This host must be a part of the Veritas InfoScale Operations Manager Management Server domain.

Table 10-27 Edit configuration panel options for IBM SVC enclosure *(continued)*

Field	Description
Enclosure IP/Name	Modify the name, or the IP address of the IBM SVC enclosure.
User Name	Modify the user name for the enclosure. Specify admin as the username for IBM SVC enclosure with version 5.1.
Password	Change the password for the enclosure. Leave this field blank if the discovery is done using the ssh private key file.
Certificate Location	Modify the full path for the ssh private key file on the discovery host.

Edit Configuration panel options to select an enclosure for deep discovery

Use this wizard panel to update the deep discovery information for an enclosure.

Table 10-28 Edit configuration options to update the deep discovery information for enclosures

Field	Description
Configuration Name	Displays the name of the deep discovery operation.
Enclosures to be configured for discovery	
Display Name	Displays name of the enclosure.
Vendor ID	Displays the ID that Veritas InfoScale Operations Manager generates for the enclosure.
Serial	Displays serial number of the enclosure.
Vendor	Displays manufacturer of the enclosure.
Model	Displays enclosure model information.
Product	Displays the type of the enclosure.

Clear the check box to disable deep discovery. The check box is selected by default. Select the check box in the top row to enable deep discovery for all the enclosures in the list.

Removing deep discovery for a storage enclosure

In the Storage Insight Add-on, you can stop performing the deep discovery for a storage enclosure by removing the deep discovery configuration. After you remove the deep discovery configuration, you cannot view the following information pertaining to the enclosures in your data center:

Table 10-29 Enclosure information that is unavailable after removing deep discovery configuration

Enclosure	Unavailable enclosure discovery information
HITACHI	LDEVs (logical devices), PDEVs (physical devices), array groups, replications, and thin pools
IBM XIV	Pool volumes, physical disks, storage pools, replications, and thin pools
EMC Symmetrix	Physical disks, devices, replications, and thin pools
EMC CLARiiON	Thin pools, thin LUNs, physical disks, RAID groups, and replications
NetApp	Physical disks, aggregates, array volumes, array ports, and replications NAS objects: Qtree, share, quota, vFiler, and consumer storage pool details
HP Enterprise Virtual Array (EVA)	Virtual disks, physical disks, array disk groups, and replications
EMC Celerra	Disk volumes, storage pools , array volumes, NAS file systems, shares, and Data Movers
IBM System Storage DS	Logical volumes, Disk Drive Modules (DDMs), Extent pools, and Ranks
EMC VNX	LUNs, physical disks, RAID groups, replications, thin pools, and information about NAS objects (for example, Shares) The removal of deep discovery information applies to all three types of VNX configurations - File, Block, and Unified.
EMC VPLEX	Virtual volumes, storage volumes, RAID groups, and replications
3PAR	Virtual volumes, physical disks, replications, Common Provisioning Groups (CPGs), and thin CPGs

Table 10-29 Enclosure information that is unavailable after removing deep discovery configuration *(continued)*

Enclosure	Unavailable enclosure discovery information
IBM SVC	Storage Pools, Virtual Disks, Managed Disks, Replications, and Adapters.

To perform this task, assign the Admin role to your user group on the Management Server perspective.

To remove deep discovery configuration from a storage enclosure

1. In the Home page on the Management Server console, click **Settings**.
2. Click **Device**.
3. Expand **Enclosures** to locate the vendor.
4. In the vendor configurations list, right-click the enclosure and select **Remove Configuration**.
5. In the **Unconfigure Device** panel, click **Yes**.

See "About Storage Insight Add-on" on page 51.

See "About the discovery host" on page 53.

Refreshing the enclosures that are configured for deep array discovery

Using the Management Server console, you can refresh the enclosures that you have configured for deep array discovery.

To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

To refresh the enclosures that are configured for deep array discovery

1. In the Home page on the Management Server console, click **Settings**.
2. Click **Device**.
3. Select **Enclosures**.
4. In the **Enclosure Configurations** tab locate the enclosure.
5. Right-click the enclosure and select **Refresh Configuration**
6. In the **Refresh Configuration** panel click **Yes**.

See "About Storage Insight Add-on" on page 51.

Monitoring the usage of thin pools

The Storage Insight Add-on lets you monitor the subscribed and the consumed size of the thin pools on your enclosures. Using the Storage Insight Add-on, you can set high and low threshold levels for the subscribed and the consumed size of the thin pools. When the consumed size of the thin pools reaches the high threshold level, Veritas InfoScale Operations Manager raises a fault alert to Management Server. The low threshold for the consumed size results in a risk alert. Similarly, when the subscribed size reaches the low or the high threshold level, Veritas InfoScale Operations Manager raises a risk or a fault alert, respectively.

By default, Veritas InfoScale Operations Manager raises a fault alert when the consumed size of the thin pool reaches 90 percentage.

The usage of thin pools can be monitored more frequently.

You can perform this operation only on the following types of enclosures:

- EMC Symmetrix
- EMC CLARiiON
- HITACHI
- IBM XIV
- HP 3PAR
- IBM SVC
- EMC VNX

Note: The values that you specify for the threshold levels come into effect only from the subsequent discoveries of the managed hosts

---

To monitor the usage of the thin pools

1. In the Management Server console, go to the **Storage** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Enclosures** to locate and select the enclosure.
3. Expand the required enclosure and select **Thin Pools**.
4. Right-click on the selected thin pool and select **Monitor Thin Pool**.

5 In the **Monitor Thin Pool** wizard panel, set the various threshold levels for monitoring the thin pools.

See "Monitor Thin Pool panel options" on page 97.

6 In the **Result** panel, verify that the operation is successful.

See "About performance metering statistics for enclosure" on page 99.

See "About Storage Insight Add-on" on page 51.

Monitor Thin Pool panel options

Use this wizard panel to set threshold levels for monitoring the usage of the subscribed and the consumed space in the thin pools on your enclosures.

Using this panel, you can set:

- Different threshold levels for the thin pools
- The same threshold levels for all the thin pools

Table 10-30 Monitor Thin Pool panel options

Field	Description
Apply to all thin pools	Select this check box to set the same threshold levels for all the thin pools that are listed on this panel.
Low Subscription(%)	Enter a value to specify the low threshold level for the subscribed size of the thin pools. When the subscribed size of the thin pools reaches this level, Veritas InfoScale Operations Manager raises a risk alert. This value is applied to all the thin pools in the wizard panel.
High Subscription(%)	Enter a value to specify the high threshold level for the subscribed size of the thin pools. When the subscribed size of the thin pools reaches this level, Veritas InfoScale Operations Manager raises a fault alert. This value is applied to all the thin pools in the wizard panel.

Table 10-30 Monitor Thin Pool panel options *(continued)*

Field	Description
Low Consumption(%)	Enter a value to specify the low threshold level for the consumed size of the thin pools. When the consumed size of the thin pools reaches this level, Veritas InfoScale Operations Manager raises a risk alert. This value is applied to all the thin pools in the wizard panel.
High Consumption(%)	Enter a value to specify the high threshold level for the consumed size of the thin pools. When the consumed size of the thin pools reaches this level, Veritas InfoScale Operations Manager raises a fault alert. By default, Veritas InfoScale Operations Manager raises a fault alert when the consumed size of the thin pool reaches 90 percentage. This value is applied to all the thin pools in the wizard panel.
Name	Displays the name of the thin pool.
Low Subscription(%)	Displays the low subscription value entered.
High Subscription(%)	Displays the high subscription value entered.
Low Consumption(%)	Displays the low consumption value entered.
High Consumption(%)	Displays the high consumption value entered.
% Subscribed	Displays the amount of the subscribed storage space in the thin pool.
% Consumed	Displays the amount of the consumed storage space in the thin pool.

See “Monitoring the usage of thin pools” on page 96.

Chapter 11

Monitoring storage array metering data

This chapter includes the following topics:

- About performance metering statistics for enclosure
- Disabling performance metering for an enclosure
- Enabling performance metering for an enclosure
- Viewing the performance graphs for an enclosure
- Viewing the performance graphs for an array port
- Viewing the performance graphs for an adapter

About performance metering statistics for enclosure

When you upgrade to Storage Insight Add-on from version 4.1 to 6.0 or later, metering of existing storage enclosures is automatically enabled if the discovery host version is 5.0 and later. When you upgrade from Storage Insight Add-on version 5.0 to 6.0 or later, and if you have EMC VNX (Block) enclosures configured using Storage Insight Add-on version 5.0, then metering for the enclosure is enabled by default. Metering state (enabled or disabled) for all other enclosures does not change when you upgrade to version 6.0 or later. Using the Management Server console, you can change the metering interval and, enable or disable performance metering for the selected enclosure.

For more information on performance metering statistics, see the *Veritas Operations Manager Management Server User Guide*.

Performance metering is enabled only for the following storage enclosures:

- EMC CLARiiON
- EMC Symmetrix
- EMC VNX (Block)
- IBM XIV
- NetApp

Note: Performance metering is enabled only for those NetApp enclosures that have Data ONTAP version 7.3 or later.

Following are the pre-requisites for enabling performance metering:

- NetApp: User should have execution permission for "perf-object-get-instances" and "fcp-adapter-list-info" APIs.
- EMC CLARiiON and EMC VNX (Block): Statistics gathering should be enabled on the enclosure.
- EMC Symmetrix: The symstat CLI should be installed.

Table 11-1 lists the log configurations for the performance charts.

Table 11-1 Log configurations

Resource	Chart name	Log configuration
Storage array - Port	IO Operations per second IO Throughput per second	30 minutes / 1 day 2 hours / 1 month 1 day / 1 year
Storage array - Adapter	IO Operations per second IO Throughput per second	30 minutes / 1 day 2 hours / 1 month 1 day / 1 year
Storage array - Enclosure	Average Read/Write Latency for Host Average Bytes Read/Written for Host (in bytes) IO Operations per second	30 minutes / 1 day 2 hours / 1 month 1 day / 1 year

See “Enabling performance metering for an enclosure” on page 101.

See "Disabling performance metering for an enclosure" on page 101.

See "Monitoring the usage of thin pools" on page 96.

Disabling performance metering for an enclosure

Use this option to disable performance metering for an enclosure.

To perform this task, your user group must be assigned the Admin role on the enclosure or the **Storage** perspective. The permission on the enclosure may be explicitly assigned or inherited from a parent Organization.

To disable performance metering for an enclosure

1. In the Management Server console, go to the **Storage** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Enclosures** to locate the enclosure.
3. Right-click the enclosure and select **Properties**.
4. Click the **Performance** tab, clear the **Enable performance metering** check box, and click **OK**.

See "Enabling performance metering for an enclosure" on page 101.

See "About Storage Insight Add-on" on page 51.

See "About the discovery host" on page 53.

Enabling performance metering for an enclosure

Use this option to enable performance metering for an enclosure.

You can enter an interval period for which you want the performance data to be collected. The interval period should be between 5 minutes and 1440 minutes and in the multiples of five.

To perform this task, your user group must be assigned the Admin role on the enclosure or the Storage perspective. The permissions on the enclosure may be explicitly assigned or inherited from a parent Organization.

To enable performance metering for an enclosure

1. In the Management Server console, go to the **Storage** perspective, and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Enclosures** to locate the enclosure.
3. Right-click the enclosure and select **Properties**.

4 Click the **Performance** tab.

5 Select **Enable performance metering**, and enter an interval period for which you want the performance data to be collected.

See "Disabling performance metering for an enclosure" on page 101.

See "About performance metering statistics for enclosure" on page 99.

Viewing the performance graphs for an enclosure

In the Management Server console, you can view the performance of an enclosure in an interactive graph. You can review the performance of an enclosure for different durations - 6 hours, 24 hours, 1 week, 1 month, 1 year, or based on live data.

You can view the performance graphs for an enclosure only if the storage provisioned to a host from the enclosure is managed by DMP or VxDMP. You can also view the performance charts if the enclosure is discovered by Storage Insight Add-on version 6.1 or later.

Average Read/Write Latency for Host and **Bytes Read/Written for Host** charts are not displayed for virtual machines having virtual initiator.

Table 11-2 lists the performance graphs for an enclosure.

Table 11-2 Performance graphs for an enclosure

Performance graph name	Description
Average Read/Write Latency for Host	Displays the average read and write latency and the trend for the selected host for the specified duration.
Bytes Read/Written for Host	Displays the bytes read and written (in KB) and the trend for the selected host for the specified duration
IO Operations per second	Displays the number of IO operations per second and the linear trend for the selected host for the specified duration. **Note:** You can view this graph only if Storage Insight Add-on version 6.1 or later is installed This graph is displayed only for EMC Symmetrix, EMC CLARiiON, EMC VNX (Block), NetApp and IBM XIV arrays. This graph cannot be rendered for **Live** data.

You can view these performance graphs for the enclosures, for which your user group has at least Guest role explicitly assigned or inherited from a parent

Organization. You can also view the graphs if your user group has at least Guest role assigned on the Storage perspective.

To view the performance graphs for an enclosure

1. In the Management Server console, go to the **Storage** perspective and select **Manage** in the left pane.
2. Expand Organization or **Uncategorized Enclosures** to locate the enclosure.
3. Do one of the following:
   - If the enclosure is EMC VNX (Block), expand **Components** to select the **Block**.
   - For other enclosures, skip to step 4.
4. Click the **Performance** tab. To change the duration, use the drop-down list.
5. Click the ellipses to select a host.

Viewing the performance graphs for an array port

In the Management Server console, you can view the performance of a Fibre Channel array port in an interactive graph. You can review the performance for different durations - 6 hours, 24 hours, 1 week, 1 month, 1 year, or based on live data.

For NetApp array, performance graph is displayed only if the port is online and local to the array.

Table 11-3 lists the performance graphs for an array port.

Table 11-3 Performance graphs for a Fibre Channel array port

Performance graph name	Description
IO Operations per second	Displays the number of input output operations per second and the trend for the selected array port for the specified duration.
IO Throughput per second	Displays the data read and written (in KB) per second and the trend for the selected array port for the specified duration.

You can view these performance graphs for the enclosures, for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the graphs if your user group has at least Guest role assigned on the Storage perspective.

To view the performance graphs for an array port

1. In the Management Server console, go to the **Storage** perspective and select **Manage** in the left pane.
2. Expand Organization or **Uncategorized Enclosures** to locate the enclosure.
3. Do one of the following.
   - If the enclosure is EMC VPLEX enclosure, expand **Nodes**, and then expand the cluster.
   - If the enclosure is a EMC VNX (Block) enclosure, expand **Components**, and then expand **Block**.
   - For other enclosures, skip to step 4.
4. Expand **Array Ports** to locate an array port.
5. Select the array port and click on the **Performance** tab. To change the duration, use the drop-down list.

See "About performance metering statistics for enclosure" on page 99.

Viewing the performance graphs for an adapter

In the Management Server console, you can view the performance of a Fibre Channel adapter in an interactive graph. You can review the performance for different durations - 6 hours, 24 hours, 1 week, 1 month, or 1 year.

Table 11-4 lists the performance graphs for an adapter.

Table 11-4 Performance graphs for a Fibre Channel adapter

Performance graph name	Description
IO Operations per second	Displays the number of input output operations per second and the trend for the selected adapter for the specified duration.
IO Throughput per second	Displays the data read and written (in KB) per second and the trend for the selected adapter for the specified duration.

You can view these performance graphs for the enclosures, for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the graphs if your user group has at least Guest role assigned on the Storage perspective.

To view the performance graphs for an adapter

1. In the Management Server console, go to the **Storage** perspective and select **Manage** in the left pane.
2. Expand Organization or **Uncategorized Enclosures** to locate the enclosure.
3. Do one of the following.
   - If the enclosure is EMC VPLEX enclosure, expand **Nodes**, and then expand the cluster.
   - If the enclosure is a EMC VNX (Block) enclosure, expand **Components**, and then expand **Block**.
   - For other enclosures, skip to step 4.
4. Expand **Adapters** to locate an adapter.
5. Select the adapter and click on the **Performance** tab. To change the duration, use the drop-down list.

See "About performance metering statistics for enclosure" on page 99.

See "About Storage Insight Add-on" on page 51.

Chapter 12

Managing LUN classifications

This chapter includes the following topics:

- About LUN classification
- Creating LUN classifications
- Modifying LUN classifications
- Deleting LUN classifications
- Modifying the order of the LUN classifications
- Refreshing the LUN classifications

About LUN classification

You can define classifications for LUNs in your data center. These LUN classifications can be used, for example, for LUN selection when migrating or provisioning storage with Storage Provisioning and Enclosure Migration Add-on.

LUN classification is available only when Storage Insight Add-on is installed and configured on Management Server.

You can classify LUNs based on one or more parameters, such as the following:

- Array (enclosure) name
- Whether or not the LUN is thin
- LUN name
- Product, for example CLARiiON
- RAID group

- RAID level of a LUN, for example, RAID5/6
- Whether or not the LUN is replicated
- Vendor

When you define classifications for LUNs, Veritas InfoScale Operations Manager stores these definitions as rules in Management Server.

In addition to defining classifications, you can assign them an order. In certain circumstances, multiple classification definitions can apply to a LUN. In such cases, Veritas InfoScale Operations Manager looks for the highest order of classification that applies and assigns it to the LUN.

See "Creating LUN classifications" on page 107.

See "Deleting LUN classifications" on page 109.

See "Modifying LUN classifications" on page 108.

See "Modifying the order of the LUN classifications" on page 110.

See "Refreshing the LUN classifications" on page 110.

Creating LUN classifications

Veritas InfoScale Operations Manager lets you classify LUNs based on one or more parameters, such as the following:

- Array (enclosure) name
- Whether or not the LUN is thin
- LUN name
- Product, for example CLARiiON
- RAID group
- RAID level of a LUN, for example, RAID5/6
- Whether or not the LUN is replicated
- Vendor

To perform this task, your user group must be assigned the Admin role on the Storage perspective.

To create LUN classifications

1. In the Management Server console, go to the **Storage** perspective and expand **Manage** in the left pane.
2. Right-click **Data Center**.

3 Click **LUN Classification > Create**.

4 In the **Create new classification for LUNS/LDEVS** panel, specify the classification as follows:

- Define the LUN selection criteria.
  Specify an attribute, a condition, and a value for each parameter. Text values are not case sensitive.
  To add a parameter, click the **Add** icon and choose whether to use an **AND** or **OR** operator for the new parameter.
- Specify a name for this LUN classification.
- Optionally, select the check box **Classify enclosure LUNs after rule creation** to have Veritas InfoScale Operations Manager apply the classification after you create it.
  For the most efficiency, Veritas recommends that you create all the LUN classifications first instead of having Veritas InfoScale Operations Manager apply each classification after you create it. You can then use **LUN Classification > Refresh Classification** option to apply all the classifications manually.

5 Click **Finish**.

You can create additional LUN classifications by repeating these steps.

6 Once LUN classifications are applied, you can view the data in the **LUN Classification** chart on the Data Center **Overview** tab. You may need to refresh the page for the data to display in the chart.

See "About LUN classification" on page 106.

Modifying LUN classifications

You can modify the classifications that were previously configured for LUNs.You can modify the name of the classification and the rules that are specified for defining the classification for a LUN.

To perform this task, your user group must be assigned the Admin role on the Storage perspective.

To modify LUN classifications

1 In the Management Server console, go to the **Storage** perspective and expand **Manage** in the left pane.

2 Right-click **Data Center**.

3 Click **LUN Classification > Edit**.

4 In the **Edit existing LUN/LDEV classification** panel, do the following:

- Select the classification to be modified.
- Optionally change the selection for **Classify enclosure LUNs after editing the rule**.
- Click **Next**.

5 Edit the existing LUN selection criteria as follows: Specify an attribute, a condition, and a value for each parameter. Text values are not case sensitive. To add a parameter, click the **Add** icon and choose whether to use an **AND** or **OR** operator for the new parameter.

Click **Finish**.

See "About LUN classification" on page 106.

Deleting LUN classifications

You can delete classifications that were previously configured for LUNs.

To perform this task, your user group must be assigned the Admin role on the Storage perspective.

To delete LUN classifications

1 In the Management Server console, go to the **Storage** perspective and expand **Manage** in the left pane.

2 Right-click **Data Center**.

3 Click **LUN Classification > Delete**.

4 In the **Delete LUN/LDEV classification** panel, do the following:

- Select the classification to be deleted.
- Optionally change the selection for **Classify enclosure LUNs after deleting the rule**.
- Click **Next**.

5 To confirm the deletion, click **Finish**.

See "About LUN classification" on page 106.

Modifying the order of the LUN classifications

You can modify the order in which the LUN classifications are applied to the LUNs. When multiple classifications are applicable to a LUN, Veritas InfoScale Operations Manager uses the classification that appears first in this order to classify the LUN.

To perform this task, your user group must be assigned the Admin role on the Storage perspective.

To modify the order of the LUN classifications

1. In the Management Server console, go to the **Storage** perspective and expand **Manage** in the left pane.
2. Right-click **Data Center**.
3. Click **LUN Classification > Reorder**.
4. In the **Reorder classification settings for LUNS/LDEVS** panel, select a classification and click **Move Up** or **Move Down** to change the order.
5. When you are satisfied with the order, click **Finish**.

See "About LUN classification" on page 106.

Refreshing the LUN classifications

When you choose the option to refresh LUN classifications, Veritas InfoScale Operations Manager applies the existing LUN classifications to the enclosure LUNs in the data center. You can use this option after you first create LUN classifications to apply the newly created classifications. You can also use this option after you make changes to the classifications. For new LUNs discovered by Veritas InfoScale Operations Manager, the existing LUN classifications are applied automatically to the new LUNS at regular intervals (within 24 hours).

To perform this task, your user group must be assigned the Admin role on the Storage perspective.

To refresh the LUN classifications

1. In the Management Server console, go to the **Storage** perspective and expand **Manage** in the left pane.
2. Right-click **Data Center**.
3. Click **LUN Classification > Refresh Classification**.
4. Confirm that you want to refresh the classifications and click **Finish**.

See "About LUN classification" on page 106.

Appendix A

Enclosure configuration prerequisites

This appendix includes the following topics:

- HITACHI enclosure configuration prerequisites
- EMC Symmetrix storage array configuration prerequisites
- IBM XIV enclosure configuration prerequisites
- NetApp storage enclosure configuration prerequisites
- EMC CLARiiON storage enclosures configuration prerequisites
- Hewlett-Packard Enterprise Virtual Array (HP EVA) configuration prerequisites
- IBM System Storage DS enclosure configuration prerequisites
- IBM SVC enclosure configuration prerequisites
- EMC Celerra enclosure configuration prerequisites
- EMC VNX storage enclosure configuration prerequisites
- EMC VPLEX storage enclosure configuration prerequisites
- 3PAR storage enclosure configuration prerequisites

HITACHI enclosure configuration prerequisites

To discover HITACHI enclosure, ensure that storage network physical connections, HiCommand server, and Storage Insight Add-on are properly configured. Storage Insight Add-on discovery host communicates with HiCommand server to access HITACHI enclosure.

For more information on supported versions of HiCommand server, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

Physical connection requirements for HITACHI enclosure

The physical connection requirements for HITACHI enclosure are as follows:

- Network connectivity between HiCommand server and Storage Insight Add-on discovery host.
- You should be able to connect from the discovery host to HiCommand server using the following URL:
  http://*HiCommand_server_address*:2001
  Where *HiCommand_server_address* is the IP address of HiCommand server, and 2001 is access port.

EMC Symmetrix storage array configuration prerequisites

To discover EMC Symmetrix storage arrays, ensure that your storage network's physical connections, device settings, and Storage Insight add-on are properly configured.

Note: You must configure an array for the discovery using the EMC Symmetrix Command Line Interface (SYMCLI).

For the latest support information, see the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

See "Physical connection requirements for EMC Symmetrix enclosure" on page 112.

See "Device setup requirements for EMC Symmetrix arrays" on page 113.

See "Veritas InfoScale Operations Manager setup requirements for in-band EMC Symmetrix storage arrays" on page 116.

See "Veritas InfoScale Operations Manager setup requirements to discover EMC Symmetrix storage arrays through remote SYMAPI servers" on page 117.

Physical connection requirements for EMC Symmetrix enclosure

The physical connection requirements for EMC Symmetrix enclosure are as follows:

- Fibre Channel connection between each Symmetrix array and the SAN fabric.

Device setup requirements for EMC Symmetrix arrays

This section lists the device setup requirements for the EMC Symmetrix arrays. It includes the configuration and verification of remote SYMAPI server.

See "EMC Solutions Enabler (SYMCLI) requirements for EMC Symmetrix arrays" on page 113.

See "Configuring the remote SYMAPI server" on page 113.

See "Verifying the configuration of a remote SYMAPI server" on page 115.

EMC Solutions Enabler (SYMCLI) requirements for EMC Symmetrix arrays

The EMC Solutions Enabler (SYMCLI) requirements are part of the device setup requirements for EMC Symmetrix arrays. To support discovery, install EMC Solutions Enabler on the discovery host. For information about determining which version of EMC Solutions Enabler to use, see the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

Configuring the remote SYMAPI server

Veritas InfoScale Operations Manager supports the discovery of EMC Symmetrix arrays with a remote SYMAPI server mechanism. This discovery method does not require in-band array connectivity to the host from which the EMC Symmetrix array is discovered.

For Veritas InfoScale Operations Manager to discover EMC Symmetrix arrays using a remote SYMAPI server, you need to configure the SYMAPI server. To configure the remote SYMAPI server in your environment, you need to perform two tasks:

- Ensure that the remote SYMAPI server daemon is running on the server where the EMC Solutions Enabler is installed.
- Ensure that the EMC Solutions Enabler on the discovery host can communicate with the remote SYMAPI server.

To ensure that the remote SYMAPI server daemon is running on the server where the EMC Solutions Enabler is installed

1. Log on with the administrative credentials to the host that you want to use as the remote SYMAPI server and which has in-band connectivity to the EMC Symmetrix array.
2. Type the following command on the host:

   ```
   stordaemon list
   ```

   An example of the daemon list appears.

   ```
   Available Daemons   ('[*]': Currently Running):
   [*]  storapid        EMC Solutions Enabler Base Daemon
        storgnsd        EMC Solutions Enabler GNS Daemeon
        storrdfd        EMC Solutions Enabler RDF Daemon
        storevntd       EMC Solutions Enabler Event Daemon
   [*]  storwatchd      EMC Solutions Enabler Watchdog Daemon
        storsrmd        EMC Solutions Enabler SRM Daemon
        storstpd        EMC Solutions Enabler STP Daemon
        storsrvd        EMC Solutions Enabler SYMAPI Server Daemon
   [*]  storsrvdInst    >>> Running Instance of storsrvd <<<
   ```

   The name for the remote SYMAPI server daemon is `storsrvd`. If you see a [*] for `storsrvd`, that means the remote SYMAPI server daemon is already running on the host. If the daemon is running, proceed to the next procedure.

   To ensure that the EMC Solutions Enabler can communicate with the remote SYMAPI server

3 Type the following commands to start the `storsrvd` daemon:

```
stordaemon start storsrvd

   Waiting for daemon to start. This may take several seconds.

stordaemon list
```

An example of the daemon list appears.

```
Available Daemons   ('[*]': Currently Running):
[*]  storapid        EMC Solutions Enabler Base Daemon
     storgnsd        EMC Solutions Enabler GNS Daemeon
     storrdfd        EMC Solutions Enabler RDF Daemon
     storevntd       EMC Solutions Enabler Event Daemon
[*]  storwatchd      EMC Solutions Enabler Watchdog Daemon
     storsrmd        EMC Solutions Enabler SRM Daemon
     storstpd        EMC Solutions Enabler STP Daemon
[*]  storsrvd        EMC Solutions Enabler SYMAPI Server Daemon
```

4 Perform steps 1 and 2 on each host in which you want to configure the remote SYMAPI server.

To ensure that the EMC Solutions Enabler can communicate with the remote SYMAPI server

1 Install EMC Solutions Enabler on the Discovery Host.

2 Change to the SYMAPI configuration directory. By default, the directory is:

- Solaris — /var/symapi/config
- Windows — %PROGRAMFILES%\EMC\SYMAPI\config

3 Modify the file "netcnfg" in the SYMAPI configuration directory of the host where the EMC Solutions Enabler is installed. Append the entry for the configured SYMAPI server(s) to the end of the file. The following is an example of adding entries for two SYMAPI servers:

```
#SYMAPI_SERVER - TCPIP node001 WWW.XXX.YYY.ZZZ 2707 -

DC1_SERVER - TCPIP ctrlhost_1 10.200.15.155 2707 -

DC2_SERVER - TCPIP ctrlhost_2 10.249.100.155 2707 -
```

Verifying the configuration of a remote SYMAPI server

Verify the remote SYMAPI server configuration before you perform the device setup requirements. Set environment variables to test if the SYMAPI server is configured correctly.

To verify the configuration of a remote SYMAPI server

1. Open an operating system console and log on to the host as root (Solaris) or as a user with administrator-level privileges (Windows).
2. Ensure that the SYMCLI commands are in your `PATH` environment.
3. Do one of the following:
   - On Solaris, run the following SYMCLI commands to set the server's environment variables:

     ```
     SYMCLI_CONNECT_TYPE=REMOTE; export SYMCLI_CONNECT_TYPE
     SYMCLI_CONNECT=DC1_SERVER; export SYMCLI_CONNECT
     symcfg list
     ```

   - On Windows, run the following SYMCLI commands to set the server's environment variables:

     ```
     set SYMCLI_CONNECT_TYPE=REMOTE
     set SYMCLI_CONNECT=DC1_SERVER
     symcfg list
     ```

4. Ensure that the arrays on different remote SYMAPI server hosts are discovered correctly.

   If you get an error in the output (instead of a list of the Symmetrix arrays), verify that your EMC Solutions Enabler is configured properly. If it is not configured properly, consult the EMC Solutions Enabler install guide for the commands. The install guide provides the detailed instructions on configuring the SYMAPI server and related commands.
5. To unset the environment variables, type the following commands:

   ```
   unset SYMCLI_CONNECT_TYPE
   unset  SYMCLI_CONNECT
   ```

Veritas InfoScale Operations Manager setup requirements for in-band EMC Symmetrix storage arrays

Veritas InfoScale Operations Manager discovers all in-band Symmetrix storage arrays with a Fibre Channel or SCSI connection to a discovery host where SYMCLI is installed. Veritas InfoScale Operations Manager also supports discovery of EMC Symmetrix storage arrays through remote SYMAPI servers.

Veritas InfoScale Operations Manager setup requirements to discover EMC Symmetrix storage arrays through remote SYMAPI servers

Veritas InfoScale Operations Manager supports the discovery of EMC Symmetrix arrays using a remote SYMAPI server mechanism. This discovery method does not require in-band array connectivity to the discovery host specified in the array configuration. However, the host on which the SYMAPI server is running must have in-band connectivity with the Symmetrix array.

For Veritas InfoScale Operations Manager to discover EMC Symmetrix arrays using a remote SYMAPI server, specify appropriate remote SYMAPI server (for example, DC1_SERVER) while configuring the Symmetrix array.

IBM XIV enclosure configuration prerequisites

To discover IBM XIV enclosure, ensure that storage network physical connections, XIV command line interface (XCLI) utility, and Storage Insight Add-on are properly configured. Veritas InfoScale Operations Manager accesses IBM XIV enclosure using XCLI utility that is installed on the discovery host.

For more information, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL).*

See "Physical connection requirements for IBM XIV enclosure" on page 117.

See "Device setup requirements for IBM XIV enclosure" on page 117.

Physical connection requirements for IBM XIV enclosure

The physical connection requirement for IBM XIV enclosure is as follows:

- Network connectivity between the discovery host of Storage Insight Add-on and IBM XIV enclosure.

Device setup requirements for IBM XIV enclosure

Storage Insight Add-on uses XCLI utility to communicate with IBM XIV enclosure. The XCLI utility should be installed on the discovery host in Veritas InfoScale Operations Manager.

Table A-1 Device setup requirements for IBM XIV enclosure

Requirement	Description
XCLI utility	Install XCLI utility on the discovery host in Veritas InfoScale Operations Manager. The default installation directory paths are as follows: ■ UNIX: `/opt/ibm/xcli/` ■ Windows: `C:\Program Files (x86)\XIV\` For the supported XCLI versions, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

NetApp storage enclosure configuration prerequisites

To discover NetApp enclosure using Storage Insight Add-on, ensure that storage network physical connections, NetApp server, and Storage Insight Add-on are properly configured.

For more information, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

See "Physical connection requirements for NetApp enclosure discovery" on page 118.

See "Device setup requirements for NetApp enclosure discovery" on page 119.

See "Host setup requirements for NetApp enclosure discovery" on page 119.

See "Requirements to support MultiStore virtual systems for NetApp enclosure" on page 119.

Physical connection requirements for NetApp enclosure discovery

The physical connection requirements for NetApp enclosure discovery are as follows:

- Network connectivity between the discovery host of Storage Insight Add-on and NetApp enclosure.
- You should be able to connect from the discovery host to NetApp server using HTTP and HTTPS connections. Use the following URLs to access the enclosure: https://*netapp_address/na_admin*

Port 443 is used for HTTPS connection.

http://*netapp_address/na_admin*

Port 80 is used for HTTP connection.

netapp_address is the IP address or NetApp array name, registered with the Domain Name System (DNS).

Device setup requirements for NetApp enclosure discovery

Setting up device includes NetApp server configuration and enabling support for MultiStore Virtual Systems on NetApp enclosure.

Host setup requirements for NetApp enclosure discovery

Configure the array with an IP address or name, and an administrator-level account with valid user name and password. Storage Insight Add-on uses these credentials to access the enclosure.

Requirements to support MultiStore virtual systems for NetApp enclosure

If MultiStore license is installed on the filer, Veritas InfoScale Operations Manager discovers MultiStore virtual system.

EMC CLARiiON storage enclosures configuration prerequisites

To discover EMC CLARiiON enclosure, ensure that network physical connections, Navisphere CLI, and Storage Insight Add-on are properly configured.

For more information, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL).*

See "Physical connection requirements for EMC CLARiiON enclosure" on page 119.

See "Device setup requirements for EMC CLARiiON enclosure" on page 120.

See "Using Password Security file for CLARiiON enclosure" on page 120.

See "Verifying NaviSecCLI communication with CLARiiON enclosure" on page 121.

Physical connection requirements for EMC CLARiiON enclosure

The physical connection requirements for EMC CLARiiON enclosure are as follows:

- Network connectivity between the enclosure and Storage Insight Add-on discovery host.

Device setup requirements for EMC CLARiiON enclosure

The device setup requirements for EMC CLARiiON enclosure include using password security file and verifying NaviSecCLI communication with CLARiiON enclosure.

See "Using Password Security file for CLARiiON enclosure" on page 120.

See "Verifying NaviSecCLI communication with CLARiiON enclosure" on page 121.

Using Password Security file for CLARiiON enclosure

To use Secure Sockets Layer for the discovery of CLARiiON array, you need to use Password Security file on the discovery host.

To put the password in a security file

1. Log on as root to the host that manages CLARiiON array.
2. Run the following NaviSecCLI command to create the security file:

   ```
   naviseccli -user userName -password passWord -scope
   userScope AddUserSecurity [-secfilepath
   secFileFolderName]
   ```

   Where:

   userName is a valid user name for the account in the array.

   passWord is the password for the *userName*.

   userScope is the scope of the user that you specified when you created the account in the array. It can be either local (userScope=1) or global (userScope=0).

   secFileFolderName is the directory where you want the security file to reside. This directory is the path that you specify when you configure an array for the discovery.
3. If some CLARiiON arrays have a different user name, repeat step 2 for each user name, specifying a different `secFileFolderName`.

   If the `secFileFolderName` was specified in step 2, use the same name when configuring the corresponding arrays.

Verifying NaviSecCLI communication with CLARiiON enclosure

Use the following procedure to verify the communication of NaviSecCLI with the CLARiiON enclosure.

To verify that the NaviSecCLI can communicate with CLARiiON arrays

1. Enter the following command at the command prompt of CLARiiON management host:

   naviseccli -h *arrayIPAddress* -user *userName* -password *passWord* -scope *userScope* networkadmin -get

   Where:

   arrayIPAddress is the IP address, the fully qualified domain name, or the name of one of the storage processors in the array.

   userName is a valid user name for the account in the array.

   passWord is the password for *userName*.

   userScope is the scope of the user that you specified when you created the account in the array. It can be Local, Global, or LDAP.

2. Review the information that displays:

   - If NaviSecCLI can communicate with the arrays, the following information is displayed:

     ```
     Storage Processor:                  SP B
     Storage Processor Network Name:     cx500-2spb
     Storage Processor IP Address:       10.100.18.18
     Storage Processor Subnet Mask:      255.255.248.0
     Storage Processor Gateway Address: 10.100.16.1
     ```

     If this command succeeds, Storage Insight Add-on uses the credentials to communicate with the array.

   - If NaviSecCLI cannot communicate with the array, the following information is displayed:

     ```
     Broken Pipe

     Valid IP Address with Feature not installed:
     naviseccli -h cx500a -user admin -password password -scope 0
     networkadmin -get
     Management Server - Feature software is not installed or the
     command may have been typed incorrectly
     ```

```
usage:
  metalun
  migrate
  connection
  mirror
  snapview
  analyzer
```

Hewlett-Packard Enterprise Virtual Array (HP EVA) configuration prerequisites

To discover Hewlett-Packard Enterprise Virtual Array (EVA), ensure that storage network physical connections, HP Command View EVA software, and Storage Insight Add-on are properly configured.

Storage Insight Add-on accesses HP EVA arrays through Command View EVA software. Storage Scripting System Utility (SSSU) should be installed on the discovery host of the Storage Insight Add-on, and it is used to access Command View EVA software. A Command View EVA software can manage multiple HP EVA arrays.

For more information, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

See "Physical connection requirements for HP EVA enclosure" on page 122.

See "Device setup requirements for HP EVA enclosure" on page 123.

See "Verifying CLI functionality for HP EVA enclosure" on page 123.

See "Using password file to access Command View EVA software" on page 125.

See "Verifying SSSU CLI communication with HP Command View EVA software" on page 124.

Physical connection requirements for HP EVA enclosure

The physical connection requirements for HP EVA enclosure are as follows:

- Fibre Channel connection between HP EVA storage array and Command View EVA software host that manages the array.
- Network connectivity between the discovery host and the Command View EVA software host. Fibre Channel connectivity is not required.

Device setup requirements for HP EVA enclosure

Complete the following tasks to meet the setup requirements for HP EVA enclosure. It includes setting up Command View EVA software and Storage System Scripting Utility (SSSU) on the discovery host in Veritas InfoScale Operations Manager domain.

Table A-2 Device setup requirements for HP EVA enclosure

Requirement	Description
Command View EVA software	Configure Command View EVA host with an IP address, and an administrator-level account with valid credentials - user name and password. If the Command View server is already configured, it can be re-used. However, make sure that the Command View server is accessible from the Veritas InfoScale Operations Manager discovery host.
HP Storage System Scripting Utility (SSSU) requirements	Install SSSU utility on the discovery host in Veritas InfoScale Operations Manager. The default installation paths are as follows: ■ UNIX: `/opt/CPQhsv/bin/sssu` ■ Windows: `C:\Program Files\Compaq\SSSU\sssu.exe` If you install SSSU utility at any other location, you must specify the location on the **Device Configuration** panel of Storage Insight Add-on when you configure HP EVA array. For more information, refer to the vendor documentation.

Verifying CLI functionality for HP EVA enclosure

To verify the functionality of CLI for HP EVA enclosure, perform the following tasks on the discovery host where SSSU is installed. Log on as root on UNIX discovery host, or as a user with administrative privileges on Windows discovery host.

To Verify CLI functionality for HP EVA enclosure on UNIX discovery host

1. On the discovery host, run SSSU utility in interactive mode by entering the following command:

   `/opt/CPQhsv/bin/sssu`

2. If the CLI is installed and functions correctly, the following message is displayed:

   SSSU for HP StorageWorks Command View EVA

   Version: 9.3.0

   Build: 071010A

To Verify CLI functionality for HP EVA enclosure on Windows discovery host

1. On the discovery host, double-click `sssu.exe` to run the SSSU utility in interactive mode.

   If the CLI is not available, navigate to the directory where the CLI is installed, and then run it.

2. If the CLI is installed and functions correctly, the following message is displayed:

   SSSU for HP StorageWorks Command View EVA

   Version: 9.3.0

   Build: 071010A

Verifying SSSU CLI communication with HP Command View EVA software

After you have installed the SSSU utility, you need to ensure that it can communicate with HP Command View EVA software.

To verify that SSSU CLI can communicate with HP Command View EVA software

1. Start SSSU utility in interactive mode.
2. For SSSU version 8.0, or later, enter the management appliance name, or IP address, user name, and password. For example:

   ```
   ./sssu "select manager Command_View_IP username=administrator
   password=adminPW "
   ```

   where `Command_View_IP` is the IP address of the HP Command View EVA server managing the HP EVA arrays. `administrator` is the user name for the HP Command View EVA account that Veritas InfoScale Operations Manager uses to manage HP EVA arrays. `adminPW` is the user-specified password to access Command View EVA administrator account.

3 If the password file is used, use the following command:

```
./sssu "select manager Command_View_IP username=administrator"
```

4 List the HP EVA arrays that Command View EVA software manages by entering the following command:

```
ls system
```

The `ls system` command output should list the arrays to be discovered.

Using password file to access Command View EVA software

The SSSU utility has the ability to create a password file. It eliminates the need to provide password when you configure array. The password file is present on the discovery host where you have installed the SSSU CLI. The password file along with user name is used to access Command View EVA software from the Storage Insight Add-on discovery host.

IBM System Storage DS enclosure configuration prerequisites

To discover IBM System Storage DS enclosure, ensure that storage network physical connections, DSCLI utility, password file (if required), and Storage Insight Add-on are properly configured. Veritas InfoScale Operations Manager accesses IBM System Storage DS enclosure using the DSCLI utility, which is installed on the discovery host.

For more information, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

See "Physical connection requirements for IBM System Storage DS enclosure" on page 125.

See "Device setup requirements for IBM System Storage DS enclosure" on page 126.

See "Using DSCLI Password Security file for System Storage DS enclosure" on page 126.

Physical connection requirements for IBM System Storage DS enclosure

The physical connection requirement for IBM System Storage DS enclosure is as follows:

- Network connectivity between IBM System Storage DS enclosure and Storage Insight Add-on discovery host.

Device setup requirements for IBM System Storage DS enclosure

Complete the following tasks to meet the setup requirements for the IBM System Storage DS enclosure. It includes setting up DSCLI utility on the discovery host in the Veritas InfoScale Operations Manager domain.

Table A-3 Device setup requirements for IBM System Storage DS enclosure

Requirement	Description
DSCLI utility	Install DSCLI utility on the discovery host in Veritas InfoScale Operations Manager. The default installation directory paths are as follows: ■ UNIX: `/opt/ibm/dscli/` ■ Windows: `C:\Program Files (x86)\IBM\` For the supported DSCLI version, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL).*

Using DSCLI Password Security file for System Storage DS enclosure

You can use DSCLI password file option to create an encrypted password file that can be used to access the enclosure. After you have created the encrypted password file, you need to specify the full path of password file on the **Device Configuration** wizard panel of Storage Insight Add-on. The password file feature is optional. Use `managepwfile` command to create the password file. For more information, refer to IBM DS8000 array documentation.

IBM SVC enclosure configuration prerequisites

To discover IBM SVC enclosure, ensure that storage network physical connections, SSH private key file (if required), and Storage Insight Add-on are properly configured. Veritas InfoScale Operations Manager accesses IBM SVC enclosure using SSH communication.

For more information, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

See "Physical connection requirements for IBM SVC enclosure" on page 127.

See "Using SSH private key file for IBM SVC enclosure" on page 127.

Physical connection requirements for IBM SVC enclosure

The physical connection requirement for IBM SVC enclosure is as follows:

- Network connectivity between IBM SVC enclosure and the discovery host.

Using SSH private key file for IBM SVC enclosure

SSH private key file is used for user authentication. You need to specify the full path of corresponding SSH private key file on the device configuration panel of Storage Insight Add-on. A valid SSH private key file is required up to version 6.2 of IBM SAN Volume Controller and IBM Storwize V7000. For version 6.3 and later, you need to either provide password, or you can use SSH private key file for user authentication.

EMC Celerra enclosure configuration prerequisites

For the discovery of EMC Celerra enclosure, ensure that storage network physical connections, Celerra Control Station, and Storage Insight Add-on are properly configured. Celerra XML API server and CLI are installed on Celerra Control Station. The discovery host of Storage Insight Add-on communicates with Celerra Control Station to perform enclosure discovery.

For more information, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*

See “Physical connection requirements for EMC Celerra enclosure” on page 127.

See “Device setup requirements for EMC Celerra enclosure” on page 128.

See “Host setup requirements for EMC Celerra enclosure” on page 128.

See “Enabling Celerra XML API” on page 128.

See “Starting XML API server on Celerra Control Station” on page 128.

See “Configuring XML API server on Celerra Control Station” on page 129.

Physical connection requirements for EMC Celerra enclosure

The physical connection requirements for EMC Celerra enclosure are as follows:

- Network connectivity between the discovery host and Control Station. Fibre Channel connectivity is not required.
- You should be able to connect from the discovery host to Control Station using the following URL:

https://*celerra_control_station_address*
Where *celerra_control_station_address* is the IP address, or the name of the Control Station.

Device setup requirements for EMC Celerra enclosure

The device setup requirements include setting up enclosure host and configuring Celerra XML API for deep discovery of Celerra enclosure.

See "Host setup requirements for EMC Celerra enclosure" on page 128.

See "Enabling Celerra XML API" on page 128.

See "Starting XML API server on Celerra Control Station" on page 128.

See "Configuring XML API server on Celerra Control Station" on page 129.

See "XML API servlet and XML API server configuration properties" on page 129.

Host setup requirements for EMC Celerra enclosure

Configure EMC Celerra enclosure host with an IP address and an administrator-level account with valid credentials.

Enabling Celerra XML API

The EMC Celerra Control Station controls storage system's components and provides access to Celerra software. Celerra XML API and Control Station are used for the discovery of Celerra enclosure. By default, XML API is in the disabled state. The XML API must be enabled for the Celerra enclosure discovery.

Starting XML API server on Celerra Control Station

By default, the XML API is disabled on Control Station. The XML API must be enabled for the discovery of Celerra enclosure by Storage Insight Add-on discovery host.

To start XML API server on Celerra Control Station

1. As a root user uncomment the following entry in `/nas/sys/nas_mcd.cfg`:

```
daemon "XML API Server"

executable "/nas/sbin/start_xml_api_server"

optional yes

canexit yes

autorestart yes

ioaccess no
```

2. Restart the NAS services with the following command:

```
# service nas start
```

   The XML API is now started.

Configuring XML API server on Celerra Control Station

The Celerra XML API is operational when you set up using default configuration settings. However, after starting XML API server, if any communication error is encountered, you can change configuration parameter to rectify the error.

To change XML API configuration properties

1. Log on to the Control Station, and open properties file in a text editor.
2. Edit properties file to change the required parameters.

   Refer to the vendor documentation for the details about parameters to change in XML API properties.
3. After you change the parameters that affect the servlet, you must restart Tomcat server on the Celerra Control Station.

   If the changed parameters affect XML API Server, restart the server also.

XML API servlet and XML API server configuration properties

Both the XML API Servlet and XML API Server share a set of configuration parameters. All parameters are located in the properties file, `$NAS_DB/sys/xml_api.conf`, which typically resolves to `/nas/sys/xml_api.conf`.

Among the list of properties, there are some debug flags. EMC has listed the XML API configuration properties for application development and not deployment. So many of the flags are set to true (debug mode).

Table A-4 XML API Servlet and XML API Server configuration properties

Property	Description
`xml.api.server.log`	The location of the XML API Server log is relative to the `$NAS_DB` (usually set to `/nas`) directory. Currently, the value is set to `log/cel_api.log`, which normally results in the log file being recorded in the `/nas/log/cel_api.log` file. The value of this property affects the XML API Server.
`xml.api.servlet.log`	The location of the XML API Servlet log is relative to the `$NAS_DB` directory. Currently the value is set to `log/webui/cel_api.log`, which normally results in the log file being recorded in the `/nas/log/webui/cel_api.log` file. The value of this property affects the Tomcat server.
`xml.api.servlet.logmask`	This property switches on and off certain servlet log profiles. The default value for the mask is zero, which means that the servlet does not log anything. Do not change this value, unless you do it temporarily at the request from EMC Support engineers and then reset it to zero. The value of this property affects the Tomcat server.
`xml.api.user.request.validation.flag`	If true, this property performs a full user request validation. In the case of improperly formatted request packets, it returns more meaningful diagnostic messages to the user. EMC suggests you set this property to false when the application is deployed since it adds to CPU and memory overhead. The value of this property affects the XML API Server.
`xml.api.enable.indications.ext`	If true, user applications can receive indications for configuration changes. Indications on task completions and statistics are always delivered (regardless of the value of the flag). The value of this property affects the XML API Server.
`xml.api.trace.apl.calls`	This property logs APL requests and responses in the XML API server log (property `xml.api.server.log`). This property is set to true; however, at the time of application deployment, it should be set to false. The value of this property affects the XML API Server.

Table A-4 XML API Servlet and XML API Server configuration properties *(continued)*

Property	Description
`xml.api.trace.apl.indications`	This property logs APL indications in the XML API server log (property `xml.api.server.log`). This property is set to true; however, at the time of application deployment, it should be set to false. The value of this property affects the XML API Server.
`xml.api.trace.user.requests`	This property records user application requests in the XML API server log (property `xml.api.server.log`). This property is set to true; however, at the time of application deployment, set it to false. The value of this property affects the XML API Server.
`xml.api.quota.poll.offset`	This property affects the exact time of the poll. By default, the tree quota cache is populated once a day. The time is specified in minutes, starting at midnight (Control Station local time) when the poll starts. The value of this property affects the XML API Server.

EMC VNX storage enclosure configuration prerequisites

Veritas InfoScale Operations Manager lets you perform deep discovery of EMC VNX (file and block) storage enclosures.

Configuration steps for VNX file are similar to EMC Celerra.

See "EMC Celerra enclosure configuration prerequisites" on page 127.

Configuration steps for VNX block are similar to EMC CLARiiON.

See "EMC CLARiiON storage enclosures configuration prerequisites" on page 119.

Using enclosure credentials for EMC VNX storage array discovery

You can use Storage Insight Add-on to perform the deep discovery of EMC VNX enclosures. Use same username and password for VNX file and VNX block enclosures. This section provides the information related to scope (global and local) for VNX file and VNX block enclosures.

- For global scope:

- For VNX block: In the **Configure Device** panel of Storage Insight Add-on, provide the value as 0 in the **Block IP/Name:Scope** field.
- For VNX file: Not applicable.

- For local scope:
  - For VNX block: In the **Configure Device** panel of Storage Insight Add-on, provide the value as 1 in the **Block IP/Name:Scope** field.
  - For VNX file: Not applicable.

EMC VPLEX storage enclosure configuration prerequisites

For the discovery of EMC VPLEX enclosure, ensure that storage network physical connections, VPLEX server, and Storage Insight Add-on are properly configured. Using HTTP and SSH communication, Storage Insight Add-on discovery host accesses VPLEX enclosure, and discovers storage volumes, virtual volumes, and devices.

For more information, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

See "Physical connection requirements for EMC VPLEX enclosure" on page 132.

See "Device setup requirements for EMC VPLEX enclosure" on page 132.

Physical connection requirements for EMC VPLEX enclosure

The physical connection requirement for EMC VPLEX enclosure is as follows:

- Network connectivity between the discovery host of Storage Insight Add-on and VPLEX enclosure.

Device setup requirements for EMC VPLEX enclosure

The device setup requirements for EMC VPLEX enclosure are as follows:

- You should be able to connect to VPLEX server from the discovery host. As an admin user, you should be able to access EMC VPLEX management console from Veritas InfoScale Operations Manager Management Server. Typical link to VPLEX console is:
  https://*vplex_ip_address*/smsflex/VPlexConsole.html
- SSH communication to VPLEX enclosure is functional.

3PAR storage enclosure configuration prerequisites

To discover 3PAR enclosure, ensure that the storage network connection between the enclosure and the discovery host, and Storage Insight Add-on are properly configured. The discovery host communicates directly with 3PAR enclosure using the IP address that is provided on the **Device Configuration** panel of Storage Insight Add-on. The discovery uses SSH communication.

For more information, refer to *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

Appendix B

Commands used by Management Server for deep discovery of enclosures

This appendix includes the following topics:

- HITACHI storage enclosure commands
- EMC Symmetrix storage enclosure commands
- IBM XIV storage enclosures commands
- NetApp storage enclosure commands
- EMC CLARiiON storage enclosure commands
- HP EVA storage enclosure commands
- IBM System Storage DS enclosure commands
- EMC Celerra storage enclosure commands
- EMC VNX (Block) storage enclosure commands
- EMC VNX (File) storage enclosure commands
- EMC VPLEX storage enclosure commands
- 3PAR storage enclosure commands
- IBM SVC storage enclosure commands

HITACHI storage enclosure commands

The Management Server accesses HiCommand server using HTTP protocol, and runs the following XML APIs to discover the storage objects of storage array:

Target: storage array

Sub targets:

- `PDEV`: To discover the physical disks of the enclosure.
- `LDEV`: To discover the logical units of the enclosure.
- `ArrayGroup`: To discover the array groups of the enclosure.
- `ReplicationInfo`: To discover the information about the logical unit replication.
- `JournalPool`: To discover the information about the journal pools.
- `ExternalPathInfo`: To discover the information about the external volume path.
- `Port`: To discover the information about the array port.
- `PortController`: To discover the information about the port controller in one, or all storage subsystems.
- `Path`: To discover the information about the LDEV mapping.
- `VolumeConnection`: To discover the mapping information of the external volumes for LDEVs.
- `JournalPoolTier`: To discover the information about dynamic storage tiering of journal pools.

See "Adding HITACHI storage enclosures for deep discovery" on page 54.

EMC Symmetrix storage enclosure commands

The following commands are used to discover the storage objects of EMC Symmetrix enclosure. The used SYMCLI version should be 7.0, or later.

- Command to find the version of the installed Solutions Enabler:
  `symcli -out XML`
- Command to find the number of last entry in the audit log of the enclosure:
  `symaudit -sid` *`arrayid`* `show -out XML`
- Command to find the number in the audit log of the enclosure based on the specified filter parameters.
  `symaudit -sid` *`arrayid`* `list -action_code` *`action_codes`*
  `-function_class` *`function_classes`* `-record_num` *`record_no`* `-out XML`

- Command to get the information about the Symmetrix configuration:
  `symcfg -sid arrayid list -out XML`
- Command to get the detailed information about the Symmetrix configuration:
  `symcfg -sid arrayid list -v -out XML`
- Command to get a brief information about the physical disks:
  `symdisk -sid arrayid list -out XML`
- Command to get the detailed information about the physical disks:
  `symdisk -sid arrayid list -v -out XML`
- Command to get the list of all Symmetrix devices:
  `symdev -sid arrayid list -all -out XML`
- Command to get the detailed information of all Symmetrix meta-head devices:
  `symdev -sid arrayid list -meta -v -out XML`
- Command to get the detailed information of Symmetrix devices specified in dev_list:
  `symdev -sid arrayid list -v -devs dev_list -out XML`
- Command to get the list of all Symmetrix devices mapped to front-end directors:
  `symdev -sid arrayid -SA all list -out XML`
- Command to get the list of all RDF groups:
  `symcfg -sid arrayid list -rdfg ALL -out XML`
- Command to get the details of a given RDF group:
  `symrdf -sid arrayid list -rdfg rdf_group_no -out XML`
- Command to get the list of all BCV sessions created on Symmetrix:
  `symmir -sid arrayid list -out XML`
- Command to get the list of all TimeFinder/Clone sessions created on Symmetrix:
  `symclone -sid arrayid list -v -out XML`
- Command to get the list of all TimeFinder/Snap sessions created on Symmetrix:
  `symsnap -sid arrayid list -v -out XML`
- Command to get brief information about thin pools in given array ID:
  `symcfg -sid arrayid list -pool -thin -detail -mb -out XML`
- Command to get the detailed information about the thin pools in given array ID.
  `symcfg -sid arrayid list -pool -thin -detail -v -mb -out XML`
- Command to get the policy association of a given storage group:
  `symfast -sid arrayid show -association -sg sg_name -out XML`
- Command to get the information on FAST policies:

```
symfast -sid arrayid list -fp -v -out XML
```

- Command to get the detailed information on FAST tiers:
  ```
  symtier -sid arrayid list -v -out XML
  ```
- Command to get the list of all directors:
  ```
  symcfg -sid arrayid list -DIR ALL -out XML
  ```
- Command to get the detailed information of all front-end directors:
  ```
  symcfg -sid arrayid list -SA ALL -v -out XML
  ```
- Command to get the detailed information of all Fibre front-end directors:
  ```
  symcfg -sid arrayid list -FA ALL -v -out XML
  ```
- Command to get the detailed information of all FICON directors:
  ```
  symcfg -sid arrayid list -EF ALL -v -out XML
  ```
- Command to list the records within the device masking VCMDB:
  ```
  symmaskdb -sid arrayid list database -out XML
  ```
- Command to list the devices assigned in the device masking VCMDB (applies WWN):
  ```
  symmaskdb -sid arrayid list devs -wwn hba_port -out XML
  ```
- Command to list the devices assigned by records in the device masking VCMDB:
  ```
  symmaskdb -sid arrayid list devs -out XML
  ```
- Command to list the device information by the initiator group:
  ```
  symaccess -sid arrayid list devinfo -out XML
  ```

See “Adding EMC Symmetrix storage enclosures for deep discovery ” on page 68.

IBM XIV storage enclosures commands

The following commands are used to discover various storage objects of IBM XIV enclosure:

- `pool_list`: Lists all or the specified storage pool.
  ```
  /opt/XIVGUI/xcli -x -y -u user_name -p password -m XIV Array IP Address pool_list
  ```
- `vol_list`: Lists all volumes, or a specific one.
  ```
  /opt/XIVGUI/xcli -x -y -u user_name -p password -m XIV Array IP Address vol_list pool=PoolName
  ```
- `module_list`: Lists the configuration of all or the specified modules.
  ```
  /opt/XIVGUI/xcli -x -y -u user_name -p password -m XIV Array IP Address module_list
  ```

- `fc_port_list`: Lists the status and configuration of the system's Fibre Channel (FC) ports.
  `/opt/XIVGUI/xcli -x -y -u user_name -p password -m XIV Array IP Address fc_port_list module=ComponentId`
- `version_get`: Prints the current version of the system.
  `/opt/XIVGUI/xcli -x -y -u user_name -p password -m XIV Array IP Address version_get`
- `config_get`: Displays the values of the configuration parameters.
  `/opt/XIVGUI/xcli -x -y -u user_name -p password -m XIV Array IP Address config_get`
- `system_capacity_list`: Lists the system's capacities (hard and soft).
  `/opt/XIVGUI/xcli -x -y -u user_name -p password -m XIV Array IP Address system_capacity_list`
- `ipinterface_list`: Lists the configuration of a specific IP interface or all IP interfaces.
  `/opt/XIVGUI/xcli -x -y -u user_name -p password -m XIV Array IP Address ipinterface_list`
- `disk_list`: Lists the special disk status.
  `/opt/XIVGUI/xcli -x -y -u user_name -p password -m XIV Array IP Address disk_list`
- `host_connectivity_list`: Lists the FC and iSCSI-level connectivity to a predefined host.
  `/opt/XIVGUI/xcli -x -y -u user_name -p password -m XIV Array IP Address host_connectivity_list`
- `mapping_list`: Lists the mapping of volumes to a specified host or cluster.
  `/opt/XIVGUI/xcli -x -y -u user_name -p password -m XIV Array IP Address mapping_list host=HostName`

The following table provides information on the user groups and their corresponding permissions for IBM XIV enclosure:

Table B-1 User groups and their corresponding permissions for IBM XIV enclosure

User group	Permission
Storage administrator	Allowed
Application administrator	Allowed

Table B-1 User groups and their corresponding permissions for IBM XIV enclosure *(continued)*

User group	Permission
Read-only users	Allowed
Technicians	Not allowed

See "Adding IBM XIV storage enclosures for deep discovery" on page 69.

NetApp storage enclosure commands

The following commands are used to discover various storage objects of NetApp enclosure. The NetApp Data ONTAP version must be 1.4, or later.

- `system-get-info`, `system-get-version`: To discover enclosure details.
- `vfiler-list-info`: To discover Vfilers.
- `disk-list-info`: To discover physical storage information.
- `fcp-adapter-list-info`, `iscsi-adapter-list-info`: To discover adapters.
- `aggr-list-info`: To discover aggregates.
- `volume-list-info`: To discover storage array volumes (FlexVol and traditional).
- `lun-list-info, lun-map-list-info`, `lun-get-serial-number`: To discover LUN details and mapping information.
- `snapmirror-get-status`, `snapshot-list-info`: To discover snapshot details.
- `qtree-list`: To discover array file systems (Qtrees).
- `cifs-share-list-iter-start`, `cifs-share-list-iter-next`: To discover CIFS shares.
- `nfs-exportfs-list-rules`: To discover NFS Shares.
- `quota-report-iter-start`, `quota-report-iter-next`: To discover Quotas.

The following command is used to discover metering and performance statistics for NetApp enclosures:

- `perf-object-get-instances`

See "Adding NetApp storage enclosures for deep discovery" on page 70.

EMC CLARiiON storage enclosure commands

The following commands are used to discover storage objects of EMC CLARiiON enclosure. The used NaviSecCLI version should be 6.29, or later. It uses `nopoll` option for commands.

- `getdisk -capacity -usercapacity`: Command to get the disk capacities of an enclosure.
- `getagent -model -rev:` Command to get the agent model and revision.
- `arrayname`: Command to get the array name.
- `getsp -ser:` Command to get the serial number of the storage processors.
- `-xml port -list -sp`: Command to get the array ports of the storage processors.
- `-xml getall -disk`: Command to get the physical disks of an enclosure.
- `getdisk -serial`: Command to get the serial number of the physical disks of an enclosure.
- `-xml getall -rg`: Command to get the RAID group of an enclosure.
- `storagegroup -list`: Command to get the storage group of an enclosure.
- `-xml getall -lun`: Command to get the logical units of an enclosure.
- `-xml metalun -list`: Command to get the meta LUNs of an enclosure.
- `-xml Snapview -listsnapshots`: Command to get the snapshots of the LUNs.
- `-xml Snapview -listclonegroup`: Command to get the snapview clones of the LUNs.
- `-xml storagepool -list`: Command to get the thin pools of an enclosure.
- `-xml thinlun -list`: Command to get the thin LUNs created from the thin pools.
- `lun -list -showOnly NonThin`: Command to get the thick LUNs created from the thin pools.
- `port -list -reads -writes -bread -bwrite -sp`: Command to get read or write on the array port.

The format of the command is as follows:

naviseccli -h *array ip* -User *username* -Password *password* -Scope *scope CLI*

If you have specified the security file, the format is as follows:

naviseccli -h *array ip* -secfilepath *path of Perm-File*

See “Adding EMC CLARiiON storage enclosures for deep discovery” on page 72.

HP EVA storage enclosure commands

The following commands are used to discover various storage objects of HP EVA storage enclosure:

- Command to verify the connectivity and version of SSSU utility:

```
./sssu_linux_x64 "select manager \"<command_view_server_name>\"
username=<cv_username> password=<cv_password>" "exit"
```

- Command to get all the HP EVA array’s visible to the Command View:

```
./sssu_linux_x64 "select manager \"<command_view_server_name>\"
username=<cv_username> password=<cv_password>" "select system
<eva_array_name>" "ls system full xml"
```

- Command to get all the Virtual disks with full details, for a selected HP EVA array:

```
./sssu_linux_x64 "select manager \"<command_view_server_name>\"
username=<cv_username> password=<cv_password>" "select system
<eva_array_name>" "ls vdisk full xml"
```

- Command to get all the controllers with full details, for a selected HP EVA array:

```
./sssu_linux_x64 "select manager \"<command_view_server_name>\"
username=<cv_username> password=<cv_password>" "select system
<eva_array_name>" "ls controller full xml"
```

- Command to get all the disk-groups with full details, for a selected HP EVA array:

```
./sssu_linux_x64 "select manager \"<command_view_server_name>\"
username=<cv_username> password=<cv_password>" "select system
<eva_array_name>" "ls disk_group full xml"
```

- Command to get all the Physical disks with full details, for a selected HP EVA array:

```
./sssu_linux_x64 "select manager \"<command_view_server_name>\"
username=<cv_username> password=<cv_password>" "select system
<eva_array_name>" "ls disk full xml"
```

- Command to get all the connected hosts with full details, for a selected HP EVA array:

```
./sssu_linux_x64 "select manager \"<command_view_server_name>\"
username=<cv_username> password=<cv_password>" "select system
<eva_array_name>" "ls host full xml"
```

- Command to get all the iSCSI controllers with full details, for a selected HP EVA array:

```
./sssu_linux_x64 "select manager \"<command_view_server_name>\"
username=<cv_username> password=<cv_password>" "select system
<eva_array_name>" "ls iscsi_controller full xml"
```

Note: Read-only user has access to the Command View server.

See "Adding HP EVA storage enclosures for deep discovery" on page 73.

IBM System Storage DS enclosure commands

The following commands are used to discover various storage objects of IBM System Storage DS enclosure:

- `ver`: The `ver` command displays the versions of the command-line interface, storage management console (HMC or SMC), and licensed machine code.

```
/opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
-passwd <password> ver -l -bnr off
```

- `lssi`: The `lssi` command displays a list of storage images in a storage complex.

```
/opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
-passwd <password> lssi -fmt default -bnr off
```

- `lshba`: The `lshba` command displays a report that lists the storage image host adapters and status information for each host adapter in the list.

```
/opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
-passwd <password> lshba -l -fmt xml -bnr off <storage_image_id>
```

- `lsddm`: The `lsddm` command displays a report that lists the disk drive modules and status information for each disk drive module in the list.

```
/opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
-passwd <password> lsddm -l -fmt xml -bnr off <storage_image_id>
```

- `lsioport`: The `lsioport` command displays a list of I/O ports on a specified storage image and optionally provides performance metrics for each I/O port that is listed.

```
/opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
-passwd <password> lsioport -l -fmt xml -bnr off -dev
<storage_image_id>
```

- `lsarray`: The `lsarray` command generates a report that displays a list of arrays in a storage image and the status information for each array in the list.

```
/opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
-passwd <password> lsarray -l -fmt xml -bnr off -dev
<storage_image_id>
```

- `lsrank`: The `lsrank` command generates a report that displays a list of defined ranks in a storage unit and the status information for each rank in the list.

  ```
  /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
  -passwd <password> lsrank -l -fmt xml -bnr off -dev
  <storage_image_id>
  ```

- `lsextpool`: The `lsextpool` generates a report that displays a list of extent pools in a storage unit and the status information on each extent pool in the list.

  ```
  /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
  -passwd <password> lsextpool -l -fmt xml -bnr off -dev
  <storage_image_id>
  ```

- `lsfbvol`: The `lsfbvol` command generates a report that displays a list of fixed block volumes in a storage unit and the status information for each volume in the list.

  ```
  /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
  -passwd <password> lsfbvol -l -fmt xml -bnr off -dev
  <storage_image_id>
  ```

  - `lsfbvol` is used for listing volumes having ‘track-space efficient’ space allocation method.

    ```
    /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
    -passwd <password> lsfbvol -sam tse -fmt xml -bnr off -dev
    <storage_image_id>
    ```

  - `lsfbvol` is used for listing volumes having ‘extent-space efficient’ space allocation method.

    ```
    /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
    -passwd <password> lsfbvol -sam ese -fmt xml -bnr off -dev
    <storage_image_id>
    ```

- `lsckdvol`: The `lsckdvol` command generates a report that displays a list of count key data (CKD) base and alias volumes in a storage unit and the status information for each volume in the list.

  ```
  /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
  -passwd <password> lsckdvol -l -fmt xml -bnr off -dev
  <storage_image_id>
  ```

  - `lsckdvol` is used for listing ckd volumes having ‘track-space efficient’ space allocation method.

```
/opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
-passwd <password> lsckdvol -sam tse -fmt xml -bnr off -dev
<storage_image_id>
```

- `lsckdvol` is used for listing ckd volumes having 'extent-space efficient' space allocation method.

  ```
  /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
  -passwd <password> lsckdvol -sam ese -fmt xml -bnr off -dev
  <storage_image_id>
  ```

- `lshostconnect`: The `lshostconnect` command displays a list of host connections for a storage image and the status information for each host connection in the list.

  ```
  /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
  -passwd <password> lshostconnect -l -fmt xml -bnr off -dev
  <storage_image_id>
  ```

- `lsvolgrp`: The `lsvolgrp` command generates a report that displays a list of volume groups in a storage unit and the status information for each volume group in the list.

  ```
  /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
  -passwd <password> lsvolgrp -fmt xml -bnr off -dev
  <storage_image_id>
  ```

- `showvolgrp`: The `showvolgrp` command generates a report that displays the detailed properties of a volume group.

  ```
  /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
  -passwd <password> showvolgrp -lunmap -bnr off -dev
  <storage_image_id> <Volume_Group_ID>
  ```

- `lssestg`: The `lssestg` generates a report that displays the space-efficient storage values for the entire storage image.

  ```
  /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
  -passwd <password> lssestg -l -fmt xml -bnr off -dev
  <storage_image_id>
  ```

- `lsflash`: The `lsflash` command generates a report that displays a list of FlashCopy relationships and the status information for each FlashCopy relationship in the list.

  ```
  /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
  -passwd <password> lsflash -l -fmt xml -bnr off -dev
  <storage_image_id> 0000-FFFF
  ```

- `lspprc`: The `lspprc` command generates a report that displays a list of remote mirror and copy volume relationships for a storage image and the status information for each remote mirror and copy volume relationship in the list.

  ```
  /opt/ibm/dscli/dscli -hmc1 <HMC1_IP_Address> -user <user_name>
  -passwd <password> lspprc -l -fullid -fmt xml -bnr off -dev
  <storage_image_id> 0000-FFFF
  ```

Access Control: The following table provides information on the user groups and their corresponding permissions for IBM System Storage DS enclosure:

Table B-2 User access control for IBM System Storage DS enclosure

User group	Permission
Administrator	Allowed
Physical Operator	Allowed
Logical Operator	Allowed
Copy Services Operator	Allowed
Monitor	Allowed
No Access	Not allowed

See "Adding IBM System Storage DS enclosures for deep discovery" on page 74.

EMC Celerra storage enclosure commands

The following XML APIs are used to form the XML queries that are required to get the information about EMC Celerra enclosures:

- `<ControlStationQueryParams/>`: Command to get the information about Celerra Control station.
- `<CelerraSystemQueryParams/>`: Command to get the detailed information about the Celerra systems.
- `<VolumeQueryParams/>`: Command to get the information about Celerra volumes.
- `<StoragePoolQueryParams/>`: Command to get the information about Celerra storage pools.
- `<FileSystemQueryParams/>`: Command to get the information about NAS file systems.
- `<CifsShareQueryParams/>`: Command to get the information about CIFS shares.

- `<CifsServerQueryParams/>`: Command to get the information about CIFS servers.
- `<NfsExportQueryParams/>`: Command to get the information about NFS export.
- `<UserQuotaQueryParams/>`: Command to get the information about user quota on the file systems.
- `<MoverQueryParams>`: Command to get the information about data movers.
- `<VdmQueryParams/>`: Command to get the information about virtual data movers.
- `<IscsiTargetQueryParams>`: Command to get the information about iSCSI targets.
- `<IscsiLunQueryParams/>`: Command to get the information about iSCSI LUNs.
- `<IscsiMaskQueryParams/>`: Command to get the information about iSCSI LUN masking.
- `<StorageSystemQueryParams storage=\"<storage ID>\"/>`: Command to get the detailed information about Celerra enclosure.
- `<ClariionDeviceQueryParams/>`: Command to get the information about back-end CLARiiON device.
- `<SymmDeviceQueryParams/>`: Command to get the information about back-end Symmetrix device.

See “Adding EMC Celerra storage enclosures for deep discovery” on page 75.

EMC VNX (Block) storage enclosure commands

The following commands are used to discover storage objects of EMC VNX (Block) storage enclosures. The used NaviSecCLI version should be 6.29, or later. It uses `nopoll` option for commands.

- `getdisk -capacity -usercapacity`: Command to get the disk capacities of enclosure.
- `getagent -model -rev:` Command to get the agent model and revision.
- `arrayname`: Command to get the array name.
- `getsp -ser:` Command to get the serial number of the storage processors.
- `-xml port -list -sp`: Command to get the array ports of the storage processors.
- `-xml getall -disk`: Command to get the physical disks of the enclosure.

- `getdisk -serial`: Command to get the serial number of physical disks of enclosure.
- `-xml getall -rg`: Command to get the RAID group of the enclosure.
- `storagegroup -list`: Command to get the storage group of the enclosure.
- `-xml getall -lun`: Command to get the logical units of the enclosure
- `-xml metalun -list`: Command to get the meta LUNs of the enclosure
- `-xml Snapview -listsnapshots`: Command to get the snapshots of the LUNs.
- `-xml Snapview -listclonegroup`: Command to get the snapview clones of the LUNs.
- `-xml storagepool -list`: Command to get the thin pools of enclosure.
- `-xml thinlun -list`: Command to get the thin LUNs created from the thin pools.
- `lun -list -showOnly NonThin`: Command to get the thick LUNs created from the thin pools.
- `port -list -reads -writes -bread -bwrite -sp`: Command to get read or write on the array port.

The format of the command is as follows:

naviseccli -h *array ip* -User *username* -Password *password* -Scope *scope CLI*

If you have specified the security file, the format is as follows:

naviseccli -h *array ip* -secfilepath *path of Perm-File*

See "Adding EMC VNX storage enclosures for deep discovery" on page 76.

EMC VNX (File) storage enclosure commands

The following XML APIs are used to form the XML queries that are required to get the information about EMC VNX (File) storage enclosure:

- `<ControlStationQueryParams/>`: Command to get the information about Celerra Control station.
- `<CelerraSystemQueryParams/>`: Command to get the detailed information about Celerra systems.
- `<VolumeQueryParams/>`: Command to get the information about Celerra volumes.
- `<StoragePoolQueryParams/>`: Command to get the information about Celerra storage pools.

- `<FileSystemQueryParams/>`: Command to get the information about NAS file systems.
- `<CifsShareQueryParams/>`: Command to get the information about CIFS shares.
- `<CifsServerQueryParams/>`: Command to get the information about CIFS servers.
- `<NfsExportQueryParams/>`: Command to get the information about NFS export.
- `<UserQuotaQueryParams/>`: Command to get the information about user quota on the file systems.
- `<MoverQueryParams>`: Command to get the information about data movers.
- `<VdmQueryParams/>`: Command to get the information about virtual data movers.
- `<IscsiTargetQueryParams>`: Command to get the information about iSCSI targets.
- `<IscsiLunQueryParams/>`: Command to get the information about iSCSI LUNs.
- `<IscsiMaskQueryParams/>`: Command to get the information about iSCSI LUN masking.
- `<StorageSystemQueryParams storage=\"<storage ID>\"/>`: Command to get the detailed information about Celerra enclosure.
- `<ClariionDeviceQueryParams/>`: Command to get the information about back-end CLARiiON device.
- `<SymmDeviceQueryParams/>`: Command to get the information about back-end Symmetrix device.

See “Adding EMC VNX storage enclosures for deep discovery” on page 76.

EMC VPLEX storage enclosure commands

To get information about VPLEX objects (storage volume, extent, device), the `cluster configdump` command is run by the Veritas InfoScale Operations Manager Management Server on the VPLEX array.

Veritas InfoScale Operations Manager Management Server makes an HTTP request to the VPLEX array. The request format is as follows:

https://*vplex IP address*:443/vplex/cluster+configdump

It gets the information about all storage entities of EMC VPLEX storage system.

See “Adding EMC VPLEX storage enclosures for deep discovery” on page 78.

3PAR storage enclosure commands

Using SSH communication, the following commands are used to discover the storage objects of a 3PAR enclosure:

- `showport`: Command to discover the array ports of the enclosure.
- `showvv -d`: Command to discover the virtual volumes of the enclosure.
- `showvv -r`: Command to get the raw capacities of the virtual volumes in the enclosure.
- `showpd`: Command to get the details of the physical disks of the enclosure.
- `showpd -i`: Command to get the extended attributes of the enclosure.
- `showvv -showcols VSize_MB -p -prov tp* -cpg`: Command to get the size of the thin virtual volumes.
- `shownet`: Command to get the IP address of the enclosure.
- `showversion -s`: Command to get the version of the 3PAR enclosure.
- `showsys`: Command to get the details of the 3PAR enclosure.

See "Adding 3PAR storage enclosures for deep discovery" on page 79.

IBM SVC storage enclosure commands

The following commands are used to discover various storage objects of IBM SVC storage enclosure

- Command to get IBM SVC array model number:
  ```
  echo SVC_PRODUCT_ID:$SVC_PRODUCT_ID
  ```
- Command to get all clusters and their detailed attributes:
  ```
  svcinfo lscluster -delim : -bytes $ID
  ```
- Command to get all nodes and their detailed attributes:
  ```
  svcinfo lsnode -delim : $ID
  ```
- Command to get all managed disk groups and their detailed attributes:
  ```
  svcinfo lsmdiskgrp -delim : -bytes $ID
  ```
- Command to get all managed disks and their detailed attributes:
  ```
  svcinfo lsmdisk -delim : -bytes $ID
  ```
- Command to get all volumes and their detailed attributes:
  ```
  svcinfo lsvdisk -delim : -bytes $ID
  ```
- Command to get all hosts:

```
svcinfo lshost -delim : $ID
```

- Command to get host to vdisk correlation:

```
svcinfo lshostvdiskmap -delim : -nohdr
```

- Command to get all internal drives:

```
svcinfo lsdrive -delim : -bytes $ID
```

- Command to get all flash copy mappings:

```
svcinfo lsfcmap -delim : $ID
```

- Command to get all remote copy mappings:

```
svcinfo lsrcrelationship -delim : $ID
```

Note: The entered user credential should have corresponding SSH key pair configured. A private key file is mandatory up to version 6.2.x.x of SVC enclosure.

See “Adding IBM SVC storage enclosures for deep discovery” on page 80.

Section 6

Storage Insight SDK Add-on 7.3

- Chapter 13. Overview of Storage Insight SDK Add-on 7.3
- Chapter 14. Managing Veritas InfoScale Operations Manager Storage Insight plug-ins

Chapter 13

Overview of Storage Insight SDK Add-on 7.3

This chapter includes the following topics:

- About Storage Insight SDK Add-on
- Array objects discovered by Storage Insight plug-in
- Operational workflow to create Storage Insight plug-in
- About discovering array information using Storage Insight SDK Add-on

About Storage Insight SDK Add-on

Veritas InfoScale Operations Manager provides deep discovery information for storage arrays from multiple vendors. This discovery is enabled through Storage Insight Add-on. For the list of supported enclosures, refer to the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

If the storage discovery for an array is not supported through Storage Insight Add-on, you can use Storage Insight SDK Add-on (Storage Insight SDK Add-on) to add support for the array in Veritas InfoScale Operations Manager. You can use Storage Insight SDK Add-on to create customized array-specific add-ons. The custom add-ons are called Storage Insight plug-ins. The Storage Insight plug-in, when installed on a Veritas InfoScale Operations Manager Management Server, enables Veritas InfoScale Operations Manager to perform the deep discovery for the array.

See “Creating a Storage Insight plug-in” on page 183.

See “Editing a Storage Insight plug-in” on page 192.

See “Testing a Storage Insight plug-in” on page 193.

Array objects discovered by Storage Insight plug-in

You can use Storage Insight plug-in to discover the following objects of an array:

- Physical disks
- Logical disks
- Physical to logical disk mapping
- RAID groups and thin pools
- Aggregate capacities at physical, RAID groups, and logical levels
- Correlation with agentless and virtual hosts. It is enabled through modeling of logical disk to host mapping information.
- Correlation with agent hosts. It is enabled through device identifier definitions.
- Support for local and remote replications.
- SAN support; NAS discovery is not supported.
- All other features that are supported by Storage Insight Add-on. For example, Veritas InfoScale Operations Manager tiering, and deep storage awareness for Storage Foundation.

Enclosure objects representation in Veritas InfoScale Operations Manager console:

The enclosure objects that are discovered by Storage Insight Plug-in are listed on the Veritas InfoScale Operations Manager console. The discovered objects are listed under the tab name that you have provided as the object terminology on Storage Insight SDK Add-on console. You can view the details related to the configured enclosure in the Storage perspective of the Veritas InfoScale Operations Manager Management Server console.

See “About Storage Insight SDK Add-on” on page 152.

See “About discovering array information using Storage Insight SDK Add-on” on page 154.

Operational workflow to create Storage Insight plug-in

Storage Insight SDK Add-on deployment and Storage Insight plug-in creation are two required processes to discover the enclosure information using Storage Insight plug-in.

Deploy Storage Insight SDK Add-on on a managed host, and use Storage Insight SDK Add-on console to create array-specific Storage Insight plug-in. Upload and install Storage Insight plug-in to Veritas InfoScale Operations Manager Management Server. After you have installed the Storage Insight plug-in successfully, use Storage Insight Add-on to discover enclosure properties.

See "About discovering array information using Storage Insight SDK Add-on" on page 154.

See "About creating Storage Insight plug-in" on page 155.

See "Creating a Storage Insight plug-in" on page 183.

See "Editing a Storage Insight plug-in" on page 192.

See "Testing a Storage Insight plug-in" on page 193.

About discovering array information using Storage Insight SDK Add-on

This process includes the following steps:

- Download Storage Insight SDK Add-on.
- Install Storage Insight SDK Add-on on a managed host.
  See " About installing Storage Insight SDK Add-on" on page 157.
- Create discovery script and additional scripts, provide enclosure capabilities that you want to discover, select the device identifiers, provide other enclosure details, and create the Storage Insight plug-in.
  See "About creating Storage Insight plug-in" on page 155.
- Upload the Storage Insight plug-in to Veritas InfoScale Operations Manager Management Server.
- Use Storage Insight Add-on to configure and discover array properties.

See "About Storage Insight SDK Add-on" on page 152.

See "Creating a Storage Insight plug-in" on page 183.

See "Editing a Storage Insight plug-in" on page 192.

See "Testing a Storage Insight plug-in" on page 193.

Chapter 14

Managing Veritas InfoScale Operations Manager Storage Insight plug-ins

This chapter includes the following topics:

- About creating Storage Insight plug-in
- Creating a Storage Insight plug-in
- Editing a Storage Insight plug-in
- Testing a Storage Insight plug-in

About creating Storage Insight plug-in

You can use Storage Insight SDK Add-on to enter the metadata about an array that you want to discover, and create array-specific plug-in. This topic lists the steps to develop the Storage Insight plug-in:

- Identify the array objects to be discovered, and their representation in Veritas InfoScale Operations Manager. For example, physical devices, logical devices, physical to logical device mapping, RAID groups, and thin pools.
  See "Array objects discovered by Storage Insight plug-in" on page 153.
- Identify the array vendor CLI that can be used to discover the array objects.

- Identify a mechanism to invoke the vendor CLI. The vendor CLI can be directly invoked from the discovery host where it is installed, or you may need to log on to the array. For example, use Secure Shell (SSH) to invoke the CLI.
- Define the mechanism to generate globally unique identifiers for the objects using the data that is provided by the array vendor CLI.
- Create discovery script. The discovery script invokes the vendor CLI to discover the array and its objects. When the discovery script is ready, you can create Storage Insight plug-in using Storage Insight SDK Add-on console. Launch Storage Insight SDK Add-on console, and select the **Create SI plug-in** option.
  See "About discovery script" on page 157.
- Along with the discovery script, you can also upload additional helper scripts.
  See "About additional scripts" on page 180.
- Create formulas in terms of SCSI pages and bytes to construct the unique identifier for the array, and logical disks (LUNs).
- From the host, launch the following URL:
  https://*SDK-hostname*:5634/admin/htdocs/sisdk.html
  Where, *SDK-hostname* is the host where you have installed Storage Insight SDK Add-on.
- Create Storage Insight plug-in: To create Storage Insight plug-in, select the **Create SI plug-in** option on Storage Insight SDK Add-on panel. When you create the Storage Insight plug-in, you need to upload the discovery script, and specify any runtime arguments required by the discovery script.
  You also need to specify the characteristics of the array that you want to discover, and how the characteristics are represented in Veritas InfoScale Operations Manager. The specified discovery script implements the discovery of the selected characteristics.
  To support end-to-end correlation of storage discovered by the agent hosts with the storage discovered out of band by Storage Insight plug-in, you also need to provide specific formulas to be used to construct the identifiers for arrays and logical disk objects. The Storage Insight plug-in is deployed as Veritas InfoScale Operations Manager add-on.

Note: A plug-in is uniquely identified by the combination of its vendor name, product, and the discovery method. Different plug-ins can be created for the same vendor and product that use different discovery mechanisms.

See "Creating a Storage Insight plug-in" on page 183.
See "About the discovery script arguments" on page 158.
See "Command example" on page 159.

See “Enclosure objects discovery order” on page 160.
See “About the enclosure discovery command output” on page 161.
See “About device identifiers” on page 181.

- Test the Storage Insight plug-in.
  See “Testing a Storage Insight plug-in” on page 193.
- Install Storage Insight plug-in. The plug-in can be installed only on Veritas InfoScale Operations Manager Management Server. You can navigate to **Deployment Management** view of Veritas InfoScale Operations Manager console, upload, and install the plug-in.
  Ensure that Storage Insight Add-on is already installed on Veritas InfoScale Operations Manager Management Server before you upload the Storage Insight plug-in. The minimum required version of Storage Insight Add-on is 5.0.
  After the Storage Insight plug-in is installed successfully, you can discover the array using Storage Insight Add-on console. The array vendor, product, and discovery mechanism are listed on the **Device Configuration** wizard of Storage Insight Add-on.
  See “About Storage Insight Add-on” on page 51.

About installing Storage Insight SDK Add-on

Storage Insight SDK Add-on can be installed only on a managed host. Navigate to **Settings** > **Deployment** on Veritas InfoScale Operations Manager console to install the add-on.

See “Creating a Storage Insight plug-in” on page 183.

See “About Storage Insight SDK Add-on” on page 152.

About discovery script

The discovery script is used by Veritas InfoScale Operations Manager to discover the information about array characteristics, and physical, and logical capabilities.

The discovery script is installed on the discovery host. In Veritas InfoScale Operations Manager, to discover the array information using the Storage Insight Add-on, you need to designate a managed host as the discovery host. The discovery host must have the vendor-specific array management tools that are used by the discovery script installed on it. You need to select this discovery host on the device configuration panel of Storage Insight Add-on.

See “About Storage Insight Add-on” on page 51.

The discovery script also requires runtime arguments. For example, array host name, IP address, user name, and password. You need to enter these arguments on the configuration parameters panel of Storage Insight SDK Add-on. As you add

the arguments, a part of command that takes the user-defined arguments, is constructed. It may also include the interpreter, if required.

Note: After you have installed the Storage Insight plug-in successfully on Veritas InfoScale Operations Manager Management Server, you can view these arguments on the wizard panel for configuring deep discovery of an enclosure using Storage Insight Add-on.

You can add only one discovery script for a Storage Insight plug-in. You can use any programming language to create the discovery script. However, it must support the defined arguments to discover a specific type of object from the array. The discovery script can be invoked using any interpreter. For example, JAVA, Perl, shell, bat, Windows PowerShell and so on. You must implement the command options for discovering each array object type in the discovery script.

You can also specify additional scripts that the discovery script can use for tasks such as performing checks, or providing information about the Storage Insight plug-in.

Note: Not all additional scripts are used directly by the discovery script. For example, a pre-defined script (`check_evn.pl`) is not used by the discovery script. It is used by Veritas InfoScale Operations Manager to check if the discovery host meets the installation requirements. Similarly, `Readme.html` is also not used by the discovery script. Apart from these scripts, you can add other scripts that can be used directly by the discovery script.

See "About additional scripts" on page 180.

See "Creating a Storage Insight plug-in" on page 183.

See "Editing a Storage Insight plug-in" on page 192.

See "Testing a Storage Insight plug-in" on page 193.

About the discovery script arguments

The discovery script run-time arguments are passed to the discovery script when it is run by the Storage Insight plug-in agentlet to discover the array objects. On the **Configuration Parameters** wizard panel, you can enter these parameters. The run time arguments form the command.

The two types of arguments are as follows:

- User-defined: These arguments are required by the underlying vendor CLI, which is called by the discovery script to access the array.

For example, the command `discover-array.pl --ip ${ip} --user ${user} --password ${password}` has three run-time arguments. They are `ip`, `user`, and `password`.

After the Storage Insight plug-in is installed successfully, you can view these arguments on the **Device Configuration** wizard panel of Storage Insight Add-on. You need to enter the values for these arguments.

- Predefined: The predefined arguments are the command line options to discover a type of array object.

  For example, `discover-array.pl --list pdevs`. In this example, `--list pdevs` is the command option that the discovery script needs to implement. The command should list the records of all physical disks discovered from the configured array. The output format of the record is predefined.

When the above-mentioned command is run, the user-defined arguments are also passed to it. In this example, the complete command line called by the agentlet to discover the physical disks is as follows:

```
"discover-array.pl --ip 10.2.1.0 --user admin --password adminpassword
--list pdevs"
```

See "About discovery script" on page 157.

See "Command example" on page 159.

See "Enclosure objects discovery order" on page 160.

Command example

This topic describes how a command is constructed using Storage Insight SDK Add-on.

- You use Perl to create the discovery script `discover-array.pl`.
- You enter username, IP address, and password as the run-time arguments, and select the following CLI options for them:
  - For username: `username`
  - For password: `pwd`
  - For IP address: `IP`

  The constructed command is as follows:

  ```
  "discover-array.pl --username ${username} --pwd ${pwd} --IP ${IP}"
  ```
- You select **Discover Physical Disks** option on the **Enclosure Capabilities** panel.

At the end of the process, the following command is created:

```
"discover-array.pl --username ${username} --pwd ${pwd} --IP ${IP}
--list pdevs"
```

The values of username, IP, and password are provided by the user during the run-time, and a complete command is constructed. For example,

```
"discover-array.pl --username admin --pwd adminpassword --IP 10.2.1.0
--list pdevs"
```

The above example provides information about the physical disk. However, you can select logical devices, array ports, adapters, replications, logical groups as the discovery object. For each option, a separate command is created.

Note: If the discovery script is a Perl script, Veritas InfoScale Operations Manager implicitly invokes it with the Perl that is provided with `VRTSsfmh` package. It is present at `sfmh-install-dir/bin/perl`.

See "Enclosure objects discovery order" on page 160.

See "About discovery script" on page 157.

See "About the enclosure discovery command output" on page 161.

Enclosure objects discovery order

The Storage Insight SDK Add-on creates an agentlet that calls the discovery script to discover the enclosure and its objects. The agentlet is bundled with the Storage Insight plug-in. The array objects are discovered in the following order:

- During new device configuration, Veritas InfoScale Operations Manager runs `check_env.pl` script. Thereafter, the discovery script (with the `-list encl` option) is run.
- `--get-data` command is run.

  It is optional to implement get-data command. It is the hook provided to the discovery script so that it can pre-fetch all the required data from the array when this command is invoked. Now the subsequent commands can use the same data instead of querying the array again.
- Discovery of enclosure objects: It happens when the deep discovery is enabled after the initial probe is successful. The enclosure objects are discovered in the following order:
  - --list adapters
  - --list ports
  - --list pdevs

- --list raidgroups
- --list rgpdevmap
- --list ldevs
- --list replications
- --list meta-ldevs
- --list thinpools
- --list rgldevmap
- --list tpldevmap
- --list tpsourceldevmap
- --list ldevpdevmap
- --list ldevhostmap
- --list capacities

- --cleanup: It is the hook provided to the discovery script to perform any cleanup.

If you have not selected capability of any of the above-mentioned options, the corresponding command is not invoked.

Note: The discovery script and `check_env.pl` script are also passed an environment variable referred to as `SI_PLUGIN_DIR`. It is the location on the discovery host where the discovery script, or the additional scripts are installed. `SI_PLUGIN_WORK_DIR` is another environment variable, and it is used by the Storage Insight plug-in as the work area for storing any data files. For example, if the get-data option is implemented, the fetched data can be stored at this location.

See "About creating Storage Insight plug-in" on page 155.

See "Command example" on page 159.

See "Creating a Storage Insight plug-in" on page 183.

See "Editing a Storage Insight plug-in" on page 192.

See "Testing a Storage Insight plug-in" on page 193.

About the enclosure discovery command output

This section provides the details of the data that each discovery command should discover, and the format in which the data should be listed.

Each command prints the records of the discovered objects. For example, `--list pdev` lists each physical disk in the enclosure as a separate record. A record displays its attributes, and the values of the attribute.

Record 1	Attribute 1 : Value 1
	Attribute 2 : Value 2
	Attribute 3 : Value 3
	Attribute N : Value N
Record 2	Attribute 1 : Value 1
	Attribute 2 : Value 2
	Attribute 3 : Value 3
	Attribute N : Value N
Record N	Attribute 1 : Value 1
	Attribute 2 : Value 2
	Attribute 3 : Value 3
	Attribute N : Value N

Each attribute with its value, must be on a different line. The attribute and its value should be separated by a colon (:). If you do not know the value for an attribute, you can specify it as "-".

The order of attributes within a record is not relevant. The record identifier attribute must be the first one to be listed for each record. For more information, refer to the command output files for Dell Compellent Storage Insight plug-in sample. It is stored on the managed host:

`/opt/VRTSsfmh/etc/SISdk/samples/compellent/compellent_data.tar.gz`

For the complete list of commands and attributes, refer to the specific command section.

See “About Storage Insight SDK Add-on” on page 152.

See “Array objects discovered by Storage Insight plug-in” on page 153.

--list encls command to discover all enclosures

The `--list encls` command should list all enclosures. Each enclosure should be listed as a record that starts with its record identifier attribute `Id`.

Table 14-1 --list encls

Attribute	Description	Data Type	Comments
Id	Globally unique identifier of the enclosure.	String	It should match with the identifier that is obtained from the SCSI pages via the enclosure identifier formula specified during the Storage Insight plug-in creation.
Name	Name for the enclosure instance.	String	
SerialNo	Serial Number of the enclosure.	String	
FirmwareVersion	Firmware version of the enclosure.	String	
Vendor	Vendor name of the enclosure.	String	For example, the vendor for the EMC Symmetrix enclosure is EMC.
Product	The name of product.	String	For example, for EMC Symmetrix array, the product name is Symmetrix.
Model	Model of the product.	String	For example, for EMC Symmetrix array, the models are DMX and VMAX.
IPAddress	The IP address of the enclosure.	String	
AutoComputeCaps	Indicates if Storage Insight plug-in agentlet should automatically compute the aggregate capacities in the waterfall chart using the objects that are discovered by other discovery commands	Boolean	The expected values are 1 or 0. Specify 1 to enable auto compute for aggregate capacities. If it is 0, it is expected that `-list capacities` command is implemented.
AltId	Optional alternate unique identifier for the enclosure. It should be globally unique.	String	It is an optional argument. It is useful while identifying a remote enclosure which is the target for the replication.

--list pdevs --encl *enclosure Id* command to discover all physical disks in an enclosure

The `--list pdevs --encl enclosure Id` command should list all physical disks in the enclosure specified by `-encl` argument. *enclosure Id* is the identifier for the enclosure as reported by the `-list encls` command.

Each physical disk should be listed as a record that starts with the record identifier attribute `Id`

Table 14-2 --list pdevs --encl *enclosure Id*

Attribute	Description	Data type	Comments
Id	Globally unique identifier for the physical disk.	String	One way to construct the identifier is by using the combination of physical disk name and enclosure identifier.
Name	Name of the physical disk.	String	
PhysicalSize	Total raw size of the physical disk.	Decimal	Specified in megabytes.
AvailableSize	Portion of the physical size that is free, or available.	Decimal	Specified in megabytes.
SpareSize	Portion of the physical size that is used as spare size (used for internal housekeeping) in the enclosure.	Decimal	Specified in megabytes.
Shared	Indicates if the physical disk is a part of multiple physical groups such as RAID groups.	Boolean	Expected value is 1 if the physical disk is shared, 0 if not shared.
RPM	RPM of the physical disk.	Integer	

Table 14-2 --list pdevs --encl *enclosure Id (continued)*

Attribute	Description	Data type	Comments
State	State of the physical disk. Indicates if the disk is healthy.	Integer	Expected value is 0 if the physical disk is healthy, and 1 if it is not healthy.
MediaType	Media type of the physical disk.	String	Expected values are one of the following: HDD and SSD
VendorMediaType	Additional classification for the media type.	String	Allowed values are SED, FC, SATA, BD, SAS, NL-SAS, and SAS (SED).

--list ldevs --encl *enclosure Id* command to discover all logical disks for an enclosure

The `--list ldevs --encl enclosure Id` command should list all logical disks in the enclosure specified by `-encl` argument. *enclosure Id* is the identifier for the enclosure as reported by `-list encls` command.

Each logical disk should be listed as a record that starts with the record identifier attribute `Id`.

Table 14-3 --list ldevs --encl *enclosure Id*

Attribute	Description	Data type	Comments
Id	Globally unique identifier for the logical disk.	String	This identifier should match the identifier obtained from the SCSI pages by logical disk identifier formula specified during the Storage Insight plug-in creation.
Name	Name of the logical disk.	String	
LogicalSize	Logical size of the logical disk.	Decimal	Specified in megabytes. For thin logical disks, it should be the subscribed size, and not the consumed size.
RawSize	Raw physical storage assigned to the logical disk.	Decimal	Specified in megabytes.

Table 14-3 --list ldevs --encl *enclosure Id (continued)*

Attribute	Description	Data type	Comments
Thin	Indicates if the logical disk is thin.	Boolean	Expected value is 1 if the logical disk is thin, and 0 for thick disk.
ThinConsumedSize	Logical consumed size of a thin logical disk.	Decimal	Specified in megabytes. Applicable only if the logical disk is thin. For a thin logical disk, this is consumed logical size, and not the subscribed size.
ThinConsumedRawSize	Raw physical storage consumed by a thin logical disk.	Decimal	Specified in megabytes. This capacity is applicable only to a thin logical disk. It is the raw physical storage that is consumed by the logical disk. The physical storage capacity that corresponds to the logical storage is represented by ThinConsumedSize.
RaidLevel	RAID level of the logical disk.	String	For example, RAID 5.
Layout	Vendor-specific layout for the logical disk.	String	
Assigned	Indicates if the logical disk is assigned to a host, or any front-end array.	Boolean	Expected value is 1 if the logical disk is assigned, and 0 if it is not assigned.
Virtual	Indicates if the logical disk is a pass through from any back-end array.	Boolean	Expected value is 1 if the logical disk is virtual, and 0 if it is not virtual.
Administrative	Indicates if the logical disk is used for administrative purposes.	Boolean	Expected value is 1 if the logical disk is administrative, and 0 if it is not administrative.

Table 14-3 --list ldevs --encl *enclosure Id (continued)*

Attribute	Description	Data type	Comments
FaultTolerant	Indicates if the logical disk is fault tolerant. That is, if it is protected by RAID.	Boolean	Expected value is 1 if the logical disk is fault tolerant, and 0 if it is not fault tolerant.
AltId	Optional alternate unique identifier for the logical device. It should be globally unique.	String	Optional argument; useful while identifying the target of a remotely replicated logical device. Refer `--list-replications` command for more information.

--list adapters --encl *enclosure Id* command to discover all adapters of an enclosure

The `--list adapters --encl` *`enclosure Id`* command should list all adapters in the enclosure specified by `-encl` argument. *enclosure Id* is the identifier for the enclosure as reported by the `-list encls` command.

Each adapter should be listed as a record that starts with the record identifier attribute `Id`.

Table 14-4 --list adapters --encl *enclosure Id*

Attribute	Description	Data type	Comments
Id	Globally unique identifier for the adapter.	String	
Name	Name of the adapter.	String	
Type	The type of adapter.	String	Expected values are FC, iSCSI, SAS and FCoE.

--list ports --encl *enclosure Id* command to discover all ports for an enclosure

The `--list ports --encl` *`enclosure Id`* command should list all the ports in the enclosure specified by `-encl` argument. *enclosure Id* is the identifier for the enclosure as reported by the `-list encls` command.

Each port should be listed as a record that starts with the record identifier attribute `Id`.

Table 14-5 --list ports --encl *enclosure Id*

Attribute	Description	Data type	Comments
Id	Globally unique identifier for the port.	String	It is typically the World Wide Port Name (WWPN) of the port.
Name	Name of the port.	String	
Type	The type of the port.	String	Allowed values are FC, iSCSI, SAS and FCoE.
WWPN	WWPN of the port.	String	
WWNN	WWNN of the port.	String	
AdapterId	Identifier of the port's adapter on the enclosure.	String	
State	State of the port.	Boolean	Expected value is 0 if the port is up, and 1 otherwise.

--list capacities --encl *enclosure Id* command to discover aggregate physical, RAID group, and logical capacities for an enclosure

The `--list capacities --encl enclosure Id` command should list aggregate physical, RAID group, and logical capacities of the enclosure specified by `-encl` argument. *enclosure Id* is the identifier for the enclosure as reported by the `-list encls` command. These capacities are displayed on the waterfall chart in the enclosure overview page of Veritas InfoScale Operations Manager console.

This command expects a single record that lists all the capacity values as described below:

Table 14-6 --list capacities --encl *enclosure Id*

Attribute	Description	Data type	Comments
Id	Globally unique identifier for the enclosure.	String	It is the record identifier.

Table 14-6 --list capacities --encl *enclosure Id (continued)*

Attribute	Description	Data type	Comments
PhySize	Total raw physical storage present in the enclosure.	Decimal	Specified in megabytes.
PhyConfigSize	Total raw physical storage in the enclosure that is configured. This is the capacity that is either allocated to some type of array group like RAID group, or thin pool, or it is directly assigned to logical disks.	Decimal	Specified in megabytes.
PhySpareSize	Total physical storage size that is used for internal housekeeping.	Decimal	Specified in megabytes.
PhyUnConfigSize	Total physical storage that is not configured. It should not include the spare capacity.	Decimal	Specified in megabytes.
RGSize	Total logical size of all RAID groups in the enclosure.	Decimal	Specified in megabytes. Applicable only if the array is RAID group capable.
RGFreeSize	Total logical size of all RAID groups in the enclosure that is free, and not yet carved out into the logical disks.	Decimal	Specified in megabytes. Applicable only if the array is RAID group capable.
RGAddrSize	Total logical addressable size of all RAID groups in the enclosure. That is, the logical size of RAID groups that is carved out into the logical disks.	Decimal	Specified in megabytes. Applicable only if the array is RAID group capable.
RGOverheadSize	Total overhead size of all RAID groups in the enclosure.	Decimal	Specified in megabytes. Applicable only if the array is RAID group capable.
LogicalSize	Total logical storage size in the enclosure.	Decimal	Specified in megabytes. Applicable only if the array is RAID group capable

Table 14-6 --list capacities --encl *enclosure Id (continued)*

Attribute	Description	Data type	Comments
LogicalAssignedSize	Total logical storage in the enclosure that is assigned to hosts, or any front-end array.	Decimal	Specified in megabytes.
LogicalUnassignedSize	Total logical storage in the enclosure that is not assigned to hosts, or any front-end array.	Decimal	Specified in megabytes.
LogicalAdminSize	Total logical storage in the enclosure that is used for the administrative purpose.	Decimal	Specified in megabytes.
LogicalRepTargetSize	Total logical storage in the enclosure that is used as replication target.	Decimal	Specified in megabytes.

See "About the enclosure discovery command output" on page 161.

See "About creating Storage Insight plug-in" on page 155.

See "About discovery script" on page 157.

--list ldevpdevmap --encl *enclosure Id* command to discover logical disk-physical disk mapping for an enclosure

The `--list ldevpdevmap --encl` *`enclosure Id`* command should list the physical disks whose storage is consumed by logical disks in the enclosure that is specified by `-encl` argument. *enclosure Id* is the identifier for enclosure as reported by the `-list encls` command.

Each record should list the logical disk, and the corresponding physical disk it gets the storage from. If a logical disk consumes storage from multiple physical disks, there are multiple records for the same logical disk. The attribute `Id` is the record identifier.

Table 14-7 --list ldevpdevmap --encl *enclosure Id*

Attribute	Description	Data type
Ldev	Globally unique identifier of the logical disk.	String

Table 14-7 --list ldevpdevmap --encl *enclosure Id (continued)*

Attribute	Description	Data type
Pdev	Globally unique identifier of the corresponding physical disk.	String

See "About the enclosure discovery command output" on page 161.

See "About creating Storage Insight plug-in" on page 155.

See "About discovery script" on page 157.

--list ldevhostmap --encl *enclosure Id* command to discover logical device-host mapping for an enclosure

The `--list ldevhostmap --encl enclosure Id` command should list the masking information for logical disks in the enclosure that is specified by `-encl` argument. *enclosure Id* is the identifier for the enclosure as reported by the `-list encls` command.

Each record should list the masking information as described below:

Table 14-8 --list ldevhostmap --encl *enclosure Id*

Attribute	Description	Data type	Comments
Id	Globally unique identifier of the logical disk.	String	
HbaPortWWPN	World Wide Port Name (WWPN) of the HBA port to which the logical disk is mapped.	String	
ArrayPortWWPN	World Wide Port Name (WWPN) of the array port to which the logical disk is mapped.	String	
LunId	SCSI LunId of the logical disk in context with the array port.	String	LunId should be listed as hexadecimal number in upper case. For example, 3C.

See "About the enclosure discovery command output" on page 161.

See "About creating Storage Insight plug-in" on page 155.

See "About discovery script" on page 157.

--list meta-ldevs --encl *enclosure Id* command to discover the mapping of meta logical disks with segment logical disks for an enclosure

The `--list meta-ldevs --encl enclosure Id` command should list the mapping of meta logical disks with the corresponding segment logical disks in the enclosure that is specified by `-encl` argument. *enclosure Id* is the identifier for the enclosure as reported by the `-list encls` command.

Each record should list this mapping information as described below:

Table 14-9 --list meta-ldevs --encl *enclosure Id*

Attribute	Description	Data type
Id	Globally unique identifier of the meta logical disk.	String
SegmentId	Globally unique identifier of the segment logical disk.	String

See "About the enclosure discovery command output" on page 161.

See "About creating Storage Insight plug-in" on page 155.

See "About discovery script" on page 157.

--list raidgroups --encl *enclosure Id* command to discover RAID groups for an enclosure

The `--list raidgroups --encl enclosure Id` command should list all RAID groups in the enclosure that is specified by `-encl` argument. *enclosure Id* is the identifier for the enclosure as reported by the `-list encls` command.

Each RAID group should be listed as a record that starts with the record identifier attribute `Id`. The RAID group information is described below:

Table 14-10 --list raidgroups --encl *enclosure Id*

Attribute	Description	Data type	Comments
Id	Globally unique identifier for the RAID group.	String	
Name	Name of the RAID group.	String	

Table 14-10 --list raidgroups --encl *enclosure Id (continued)*

Attribute	Description	Data type	Comments
State	State of the RAID group.	Integer	Its value is zero if the RAID group is healthy, and 1 otherwise.
PdevCount	Total number of physical disks that create the RAID group.	Integer	
RawSize	Total raw physical size of the RAID group.	Decimal	Specified in megabytes.
OverheadSize	RAID overhead that is introduced by the RAID group.	Decimal	Specified in megabytes.
LogicalSize	Total logical size of the RAID group.	Decimal	Specified in megabytes. LogicalSize = RawSize - OverheadSize
LogicalUsedSize	Total logical size that is already carved out as logical disks.	Decimal	Specified in megabytes.
LogicalFreeSize	Total Logical size that is not yet carved out as logical disks.	Decimal	Specified in megabytes. LogicalFreeSize = LogicalSize – LogicalUsedSize
RaidLevel	RAID level of the RAID group.	String	For example, RAID 5 (6D+1P) and RAID 0.

See "About the enclosure discovery command output" on page 161.

See "About creating Storage Insight plug-in" on page 155.

See "About discovery script" on page 157.

--list thinpools --encl *enclosure Id* command to discover all thin pools for an enclosure

The `--list thinpools --encl` *`enclosure Id`* command should list all thin pools in the enclosure that is specified by `-encl` argument. *enclosure Id* is the identifier for the enclosure as reported by the `-list encls` command.

Each thin pool should be listed as a record that starts with the record identifier attribute `Id`. The thin pool information is described below:

Table 14-11 --list thinpools --encl *enclosure Id*

Attribute	Description	Data type	Comments
Id	Globally unique identifier for the thin pool.	String	
Name	Name of the thin pool.	String	
State	State of the thin pool.	Integer	Its value is zero if the thin pool is healthy, and 1 otherwise.
SourceDevCount	Total number of source physical disks or logical disks that create the thin pool. These devices are the source of storage capacity of thin pool, and not the actual logical disks carved out of the thin pool.	Integer	
VolumeCount	Total number of thin logical disks that are carved out of the thin pool.	Integer	
RawSize	Total raw physical size of the thin pool.	Decimal	Specified in megabytes.
OverheadSize	RAID overhead that is introduced by the thin pool.	Decimal	Specified in megabytes.
LogicalSize	Total logical size of the thin pool.	Decimal	Specified in megabytes. LogicalSize = RawSize - OverheadSize
LogicalUsedSize	Total logical size that is already carved out as logical disks.	Decimal	Specified in megabytes.
LogicalFreeSize	Total Logical size that is not yet carved out as logical disks.	Decimal	Specified in megabytes. LogicalFreeSize = LogicalSize – LogicalUsedSize

Table 14-11 --list thinpools --encl *enclosure Id (continued)*

Attribute	Description	Data type	Comments
RaidLevel	RAID level of the thin pool if applicable.	String	For example, RAID 5 (6D+1P) and RAID 0.
SubscribedSize	Subscribed size of the thin pool.	Decimal	Total subscribed logical size of the thin pool. It is the sum of subscribed sizes of all thin logical disks carved out of the thin pool.

See "About the enclosure discovery command output" on page 161.

See "About creating Storage Insight plug-in" on page 155.

See "About discovery script" on page 157.

--list rgpdevmap --encl *enclosure Id* command to discover RAID group-physical disk mapping for an enclosure

The `--list rgpdevmap --encl` *`enclosure Id`* command should list the mapping of RAID groups and the physical disks that constitute the RAID group in the enclosure that is specified by `-encl` argument. *enclosure Id* is the identifier for the enclosure as reported by the `-list encls` command.

Each record lists the mapping information as described below:

Table 14-12 --list rgpdevmap --encl *enclosure Id*

Attribute	Description	Data type
Rg	Globally unique identifier for the RAID group.	String
Pdev	Globally unique identifier for the physical disk.	String

See "About the enclosure discovery command output" on page 161.

See "About creating Storage Insight plug-in" on page 155.

See "About discovery script" on page 157.

--list rgldevmap --encl *enclosure Id* command to discover RAID group-logical device mapping for an enclosure

The `--list rgldevmap --encl` *`enclosure Id`* command should list the mapping of RAID groups and the logical disks that are carved out of the RAID group in the enclosure that is specified by `-encl` argument.

enclosure Id is the identifier for the enclosure as reported by the `-list encls` command. Each record lists the mapping information as described below:

Table 14-13 --list rgldevmap --encl *enclosure Id*

Attribute	Description	Data type
Rg	Globally unique identifier for the RAID group.	String
Ldev	Globally unique identifier for the logical disk.	String

See "About the enclosure discovery command output" on page 161.

See "About creating Storage Insight plug-in" on page 155.

See "About discovery script" on page 157.

--list tpsrcldevmap --encl *enclosure Id* command to discover thin pool-source logical device mapping for an enclosure

The `--list tpsrcldevmap --encl enclosure Id` command should list the mapping of thin pools, and the source logical disks that make up the thin pool in the enclosure that is specified by `-encl` argument.

enclosure Id is the identifier for the enclosure as reported by the `-list encls` command.

Note: The logical disks are the source of storage capacity of thin pool, and not the actual logical disks carved out of the thin pool. This command is applicable to arrays whose thin pools are backed by thick logical disks (not physical disks).

Each record lists the mapping information as described below:

Table 14-14 --list tpsrcldevmap --encl *enclosure Id*

Attribute	Description	Data type
Tp	Globally unique identifier of the thin pool.	String
Ldev	Globally unique identifier of the source logical disk.	String

See "About the enclosure discovery command output" on page 161.

See "About creating Storage Insight plug-in" on page 155.

See "About discovery script" on page 157.

--list tpldevmap --encl *enclosure Id* command to discover thin pool logical device mapping for an enclosure

The `--list tpldevmap --encl enclosure Id` command should list the mapping of thin pools, and the source disks that are carved out of the thin pool in the enclosure that is specified by `-encl` argument.

enclosure Id is the identifier for the enclosure as reported by the `-list encls` command.

Each record lists the mapping information as described below:

Table 14-15 --list tpldevmap --encl *enclosure Id*

Attribute	Description	Data type
Tp	Globally unique identifier of the thin pool.	String
Ldev	Globally unique identifier of the thin logical disk.	String

See "About the enclosure discovery command output" on page 161.

See "About creating Storage Insight plug-in" on page 155.

See "About discovery script" on page 157.

--list replications --encl *enclosure Id* command to discover the replications for an enclosure

The `--list replications --encl enclosure Id` command should list the information about replications in the enclosure that is specified by `-encl` argument.

enclosure Id is the identifier for the enclosure as reported by the `-list encls` command.

Each replication should be listed as a record that starts with record identifier attribute `Id`. For a replication source-target pair, single record should be reported.

Each record lists the mapping information as described below:

Table 14-16 --list replications --encl *enclosure Id*

Attribute	Description	Data type	Comments
Id	Globally unique identifier of the replication source or target logical disk.	String	

Table 14-16 --list replications --encl *enclosure Id (continued)*

Attribute	Description	Data type	Comments
Designation	It indicates if the logical disk that is specified by `Id` attribute is the source or target of the replication.	Integer	Its value is 0 if the `Id` is replication source, and 1 if it is the replication target.
PartnerId	Globally unique identifier for the replication target if `Id` refers to the replication source. And, it is the globally unique identifier for replication source, if `Id` refers to the replication target.	String	
PartnerName	Name of the logical disk that corresponds to `PartnerId`.	String	
AcrossEncls	Specifies if the replication is across enclosures. That is, either source, or target is present in a different enclosure.	Boolean	Its value is 0 if the replication is local, and 1 if the replication is remote.
RemoteEnclId	Globally unique identifier for the remote enclosure. It is applicable only to remote replication.	String	RemoteEnclId is optional if all three attributes `RemoteEnclName`, `RemoteEnclVendor` and `RemoteEnclModel` are specified.
RemoteEnclName	Name of the remote enclosure. It is applicable only to remote replication.	String	
RemoteEnclVendor	Vendor name of the remote enclosure. It is applicable only to remote replication.	String	If replication is remote and remote enclosure's vendor is not specified, it is assumed to be the same as that of the local enclosure (referred by `enclosure Id`).

Table 14-16 --list replications --encl *enclosure Id (continued)*

Attribute	Description	Data type	Comments
RemoteEnclModel	Product name of the remote enclosure. It is applicable only to remote replication.	String	If replication is remote and remote enclosure's product is not specified, it is assumed to be the same as that of the local enclosure. (referred by `enclosure Id`).
RepMode	Mode of replication – Sync or Async.	String	Expected values are Sync for synchronous replication, and Async for asynchronous replication.
RepType	Type of replication.	String	For example, Snapshot, Clone, or Mirror.
RepState	State of replication.	String	Expected values are: Split, Synchronized, CopyOnWrite, Copied, Inactive, ReverseSynchronizing, ReverseOutofSync, Synchronizing, OutofSync, Consistent, Active, Suspended, FailedOver, Initialization, Unsynchronized, Partitioned, PreCopy, CopyInProgress, CopyOnAccess
AccessMode	Access mode of the replication target.	String	Expected values are: Rdonly: If the replication target access mode is read-only. Rdwrite: If the replication target access mode is read-write.

Table 14-16 --list replications --encl *enclosure Id (continued)*

Attribute	Description	Data type	Comments
PartnerAltId	Alternate globally unique identifier for the replication target logical device. It should match the AltId field reported for the target logical device.	String	Optional. It is useful when neither `PartnerId` nor `PartnerName` are known while discovering the replications from the source enclosure. If either `PartnerId` or `PartnerName` are specified, this field should not be specified.
RemoteEnclAltId	Alternate globally unique identifier for the remote enclosure. It should match the AltId property specified for the target enclosure.	String	Optional. It is useful when neither `RemoteEnclName` nor `RemoteEnclId` are known while discovering the replications from the source enclosure. If either `RemoteEnclName` and `RemoteEnclId` are specified, this field should not be specified.

About additional scripts

You can add multiple additional scripts to be used for the Storage Insight plug-in. An additional script may be used by Veritas InfoScale Operations Manager, or the discovery script. Two types of predefined additional scripts are listed below. They are used by Veritas InfoScale Operations Manager.

- `check_env.pl` script: This script should perform specific checks on the discovery host before the discovery script is installed on it. For example, it checks if the required interpreter is installed on the discovery host, or the correct version of vendor CLI is used. This script must be created using the Perl scripting language, and it must be invoked using `VRTSsfmh` Perl. The script name must be `check_env.pl`.

  These checks are performed on the discovery host when you configure a new device using the new device Configuration Wizard panel. The `check_env.pl` script should exit with 0 return code if all validations on the discovery host are successful. It should exit with a non-zero code if any validation fails.

- `Readme.html`: Provides the information about array vendor, the array product that Storage Insight plug-in discovers, discovery mechanism, and the vendor CLI that is required by the discovery script. It also provides information about the discovery host platform, discovered array capabilities, and the capabilities that are not discovered by Storage Insight plug-in.

See "About creating Storage Insight plug-in" on page 155.

See "Array objects discovered by Storage Insight plug-in" on page 153.

See "About discovery script" on page 157.

About device identifiers

Creating unique device identifier for arrays and logical disks is required to enable end-to-end correlation of LUNs (as seen by the hosts) to the storage arrays. The device identifier creation involves specifying vendor SCSI pages, bytes to be read, and the formulas.

Device identifiers formulas

The following table lists sample values that are provided for SCSI pages and bytes. You need to enter the values on **Device Identifier** panel of Storage Insight SDK Add-on. By default, only one field is listed on this page, and no field values are selected. Additional fields can be added. You can add up to four fields.

Field Name	SCSI Page No.	Start Byte	No of Bytes(s)	Page 83 Id Type	Field Format
Field1	0x80	0	8	Not applicable since SCSI page number is selected as 0x80.	Binary
Field2	0x80	0	16	Not applicable since SCSI page number is selected as 0x80.	Binary

Enclosure identifier formula: As per the above example, for `Field 1`, on the SCSI page 0x80, first eight bytes are read in Binary format. The formula is as follows:

```
%s${VENDOR}_%s${PRODUCT}_%s${Field1}
```

The enclosure identifier formula is constructed using the `vendor`, `product`, and `Field1`, which is the first 8 bytes from the SCSI page `0x80`. The format specifier `%s` specifies that the field following the format specifier is obtained by applying the specifier to it.

Note: It is not the definitive formula to generate the enclosure ID; it is one of the possible formulas to generate the enclosure ID.

See "Command example" on page 159.

See "About discovery script" on page 157.

See "Creating a Storage Insight plug-in" on page 183.

See "Editing a Storage Insight plug-in" on page 192.

See "Testing a Storage Insight plug-in" on page 193.

Storage Insight Plug-in sample

Sample configuration files for Dell Compellent enclosure are provided with Storage Insight SDK Add-on. When you install Storage Insight SDK Add-on, a `samples` directory is created. It is located at `VRTSsfmh/etc/SISdk/samples/compellent`.

The `samples` directory includes Storage Insight plug-in file, additional scripts, readme file, and the discovery script for Dell Compellent enclosure. These files are listed below:

- `SI_DELL_COMPELLENT_POWERSHELL-1.0.0.sfa` is the Storage Insight plug-in file for Dell Compellent enclosure.
- `dellcompcli.psl` is discovery script. It uses Windows PowerShell.
- `compellent_data.tar.gz` is the tar of output of discovery commands. It can be used to run the test option with the collected data.
- `check_env.pl` is additional script.
- `Readme.html` is readme file that contains information about the Storage Insight plug-in.

Note: You can also edit `SI_DELL_COMPELLENT_POWERSHELL-1.0.0.sfa` file to see the sample input data for various wizard panels on Storage Insight SDK Add-on console. It can be a starting point for the development of a new Storage Insight plug-in.

See "About Storage Insight SDK Add-on" on page 152.

See "Creating a Storage Insight plug-in" on page 183.

See "Editing a Storage Insight plug-in" on page 192.

See "Testing a Storage Insight plug-in" on page 193.

Creating a Storage Insight plug-in

You can use Storage Insight SDK Add-on to create array plug-ins for the arrays that are currently not supported by Storage Insight Add-on.

To create the Storage Insight plug-in

1. Log on to the host, and launch Storage Insight SDK Add-on console by typing the following URL in the Web browser:

   https://<SDK-hostname>:5634/admin/htdocs/sisdk.html

2. Select the **Create SI Plug-in** option, and click **Next**.

   See "Create, edit, and test Storage Insight plug-in panel options" on page 183.

3. Enter the plug-in related details in the **Plug-in Vitals** panel, and click **Next**.

   See "Storage Insight Plug-in Vitals panel options" on page 184.

4. Enter the configuration parameters in the **Configuration Parameters** panel, and click **Next**.

   See "Configuration Parameters panel options" on page 185.

5. Enter physical and logical capabilities that you want to discover in the **Enclosure Capabilities** panel, and click **Next**.

   See "Enclosure Capabilities panel options" on page 186.

6. Enter device identifier-related information in the **Device Identifier** panel, and click **Next**.

   See "Device Identifier panel options" on page 190.

7. Review the information that you have submitted to create the plug-in on the **Confirmation** panel, and click **Create Add-on**.

   See "Confirmation panel options" on page 191.

See "About Storage Insight SDK Add-on" on page 152.

See "About creating Storage Insight plug-in" on page 155.

Create, edit, and test Storage Insight plug-in panel options

Use this wizard panel to create a new Storage Insight plug-in. You can also edit, and test an existing plug-in using this wizard panel.

See "Creating a Storage Insight plug-in" on page 183.

See "Editing a Storage Insight plug-in" on page 192.

See "Testing a Storage Insight plug-in" on page 193.

Storage Insight Plug-in Vitals panel options

Use this wizard panel to add the information about a new Storage Insight Plug-in, or edit the parameters of an existing Storage Insight Plug-in.

Table 14-17 Plug-in Vitals panel options

Field	Description
SI Plug-in Summary	
SI Plug-in Name	Enter or update the name of the Storage Insight plug-in. The name is displayed in the **Deployment Management** view of the Veritas InfoScale Operations Manager console after the Storage Insight plug-in is uploaded.
SI Plug-in Version	Enter or update the version of the Storage Insight plug-in. The supported formats for the Storage Insight Plug-in versions are x.x, x.x.x, x.x.x.x. For example, 1.0, and 1.0.1.
Enclosure Vendor	Enter or update the enclosure vendor name. For example, for EMC Symmetrix, the vendor is EMC.
Enclosure Product	Enter or update the enclosure product name. This is the array product name from the vendor. For example, Symmetrix.
Discovery Method	Enter or update the discovery method that is used by the discovery script for the enclosure discovery. For example, SymCLI for EMC Symmetrix and Navisphere for EMC CLARiiON. The discovery method name cannot contain spaces.
Upload Discovery scripts and additional scripts	
Discovery Script	Upload the discovery script for the Storage Insight plug-in. You can upload only one discovery script per plug-in.
Additional Scripts	Upload one or more additional scripts. It is optional.

Table 14-17 Plug-in Vitals panel options *(continued)*

Field	Description
Scripts details	
File name	Displays the name of the uploaded discovery script and additional scripts.
File type	Displays the file type of the scripts.

Note: The combination of these three parameters (vendor, product, and discovery method) uniquely identify a Storage Insight plug-in. Different plug-ins can be created for the same vendor and product that use different discovery methods.

See "Creating a Storage Insight plug-in" on page 183.

See "Editing a Storage Insight plug-in" on page 192.

Configuration Parameters panel options

Use this wizard panel to provide the user-defined arguments that are required by the discovery script to communicate with the array. For example, array IP address and host name. The agentlet calls the discovery script, and passes the user-defined argument values to the discovery script. The CLI argument name used for each value is defined by the user.

For **Edit** operation, you can modify the existing parameters, and create another version of Storage Insight Plug-in.

Table 14-18 Configuration Parameters panel options

Field	Description
No Configurable Parameters	Select this checkbox to disable the option to provide the configurable parameters.
Display Name	Enter or update the display name for the parameter. After the Storage Insight plug-in is installed successfully on Veritas InfoScale Operations Manager Management Server, this field is displayed on the **Device Configuration** panel of the Storage Insight Add-on. Since this parameter is passed to the CLI, you also need to provide the corresponding CLI option for the parameter.

Table 14-18 Configuration Parameters panel options *(continued)*

Field	Description
CLI Option	Enter or update the CLI option name for the parameter. The parameter is passed to the discovery script using this CLI option.
Default value	Enter or update the default value for the parameter. It is optional.
Data Type	Select the data type for the parameter. For example, String, Integer, Password, and Boolean.
Mandatory	Specify if the parameter is mandatory.
Command line to be executed for device discovery	As you add configurable parameters, the command is constructed dynamically. The user-defined part of the discovery command line is constructed automatically. It is appended with the predefined command option for discovering a particular object type when the discovery script is run by the agentlet, which is created by the Storage Insight plug-in. For discovering each object type the discovery script is called with a different command option. But the user-defined part of the command remains the same. **Note:** You can also edit the command. For example, you can add the interpreter to the command, if required.

See "Creating a Storage Insight plug-in" on page 183.

See "Editing a Storage Insight plug-in" on page 192.

Enclosure Capabilities panel options

Use this wizard panel to add or update the physical and logical capabilities of the array that you want to discover using the Storage Insight Plug-in.

Table 14-19 Enclosure Capabilities panel options

Field	Description
Physical Characteristics	
Discover Physical Disks	Select the check box to discover the physical disk of the enclosure.
Terminology for Physical Disks	Enter or update the terminology that you want to use for the discovered physical disks. By default, its value is set to Physical Disk. Upon successful discovery of an array, the discovered physical disks are listed under **Physical Disk** tab on Veritas InfoScale Operations Manager console. This field is available only when you select the **Discover Physical Disks** check box.
Discover Array Ports	Select the check box to discover the array ports of the enclosure.
Discover Adapters	Select the check box to discover the adapters of the enclosure.
Discover RAID Groups	Select the check box to discover the RAID groups of the enclosure.
Terminology	Enter or update the terminology that you want to use for the discovered RAID groups. By default, its value is set to Raid Group. The discovered RAID groups are listed under the **Raid Group** tab on Veritas InfoScale Operations Manager console. This field is available only when you select the **Discover RAID Groups** check box. Select the **Discover raid group to source physical disk mapping**check box to discover RAID group to source physical disk mapping.
Logical Characteristics	
Discover Logical Disks	Select the check box to discover the logical disks (LUNs) of the enclosure.

Table 14-19 Enclosure Capabilities panel options *(continued)*

Field	Description
Terminology for Logical Disks	Enter or update the terminology that you want to use for the discovered logical disks. By default, its value is set to Logical Disk. The discovered logical disks are listed under **Logical Disk** tab on the Veritas InfoScale Operations Manager console. This field is available only when you select the **Discover Logical Disks** check box.
Discover Logical Disk to Host Mapping	Select the check box to discover the logical disk to host mapping for the enclosure. This field is available only when you select the **Discover Logical Disks** check box. This option allows the storage discovered by the agentless hosts to be co-related with the storage discovered out of band by the Storage Insight plug-in. It also enables the correlation with the back-end enclosures if the array being discovered is a virtualizer.
Discover Logical to Physical Disk Mapping	Select the check box to discover the logical disk to physical disk mapping for the enclosure. This field is available only when you select the **Discover Logical Disks** check box.
Discover RAID group to logical disk mapping	Select the check box to discover RAID group to logical disk mapping. This field is available only when you select the **Discover Logical Disks** check box.
Discover meta logical disks	Select this check box to discover the meta logical disks. This field is available only when you select the **Discover Logical Disks** check box.

Table 14-19 Enclosure Capabilities panel options *(continued)*

Field	Description
Terminology for meta logical disks	Enter or update the terminology that you want to use for the discovered meta logical disks. By default, its value is set to Meta Head. The discovered meta logical disks are listed under this tab on the Veritas InfoScale Operations Manager console. This field is available only when you select the **Discover meta logical disks** check box.
Discover Replications	Select the check box to discover the replications for the enclosure. Select the **Discover remote replication** check box to discover remote replications. This field is available only when you select **Discover Replications** check box.
Discover thin pools	Select this check box to discover thin pools from the enclosure.
Terminology	Enter or update the terminology that you want to use for the discovered thin pools. By default, its value is set to Thin Pool. The discovered thin pools are listed under the **Thin Pool** tab on the Veritas InfoScale Operations Manager console.
Discover thin pool to source device mapping	Select this option to discover the thin pool to source device mapping. This option is available only when you select the **Discover thin pools** check box. When you select the **Discover thin pool to source device mapping** check box, additionally **Device to back up the thin pools** option is also available to you. Select the type of devices to back up the thin pools - logical or physical.
Discover thin pool to logical disk mapping	Select this option to discover thin pool to logical disk mapping. This option is available only when you select the **Discover thin pools** check box.

See “Creating a Storage Insight plug-in” on page 183.

See “Editing a Storage Insight plug-in” on page 192.

Device Identifier panel options

Use this panel to add or edit the formulas to construct the enclosure identifier and logical disk identifier. It is required as it enables the agent to perform end-to-end correlation to the logical disk and the storage array. This panel lets you specify how the enclosure formula can be constructed from the fields, which are pulled out from the specified SCSI pages.

Table 14-20 Device Identifier panel options

Field	Description
No Device Identifier Definition(s)	Select this check box to disable the option to enter the device identifier definitions.
SCSI Vendor Id	Enter or update the SCSI vendor ID.
SCSI Product Id	Enter or update the SCSI product ID.
Field Name	Displays the name of the SCSI field.
SCSI Page No	Select the SCSI page number to be read from the drop-down list.
Start Byte	Enter or update the start byte on the SCSI page to be read.
No of Byte(s)	Enter or update the number of bytes to be read on the SCSI page. The number of bytes beginning from the start byte from the SCSI page defines the field data.
Page 83 Id Type	For SCSI page 83 ID, you can specify the Identifier ID that represents the logical disk identifier instead of explicitly specifying start byte and the number of bytes. To identify the formulas, you need to refer to the SCSI specifications of the array. The SCSI specification defines which SCSI page and bytes can be used to uniquely identify enclosure, and logical disks.
Field Format	Specify the format you want to read this field in. For example, Binary, Hex, or Integer.

Table 14-20 Device Identifier panel options *(continued)*

Field	Description
Enclosure Identifier Formula	Enter or update the formula to construct the enclosure identifier using one or more fields.
Logical Disk Identifier Formula	Enter or update the formula to construct the logical disk identifier using one or more fields.

See "Creating a Storage Insight plug-in" on page 183.

See "Editing a Storage Insight plug-in" on page 192.

See "About device identifiers" on page 181.

Confirmation panel options

Use this wizard panel to review the parameters, and physical and logical capabilities of the enclosure that you want to discover using the Storage Insight plug-in.

Table 14-21 Confirmation panel options

Field	Description
Plug-in Summary	
SI Plug-in Name:	Displays the Storage Insight plug-in name.
SI Plug-in Version:	Displays the Storage Insight plug-in version.
Enclosure Vendor:	Displays the enclosure vendor for which the Storage Insight plug-in is created.
Enclosure Product:	Displays the enclosure product for which the Storage Insight plug-in is created
Discovery Method:	Displays the discovery method that is used to discover the array information.
Physical Characteristics	Displays the physical characteristics that you have chosen to discover. For example, physical disk, logical disks, array port, and adapters. The displayed command is the exact command line that the agentlet uses to discover the array object type.

Table 14-21 Confirmation panel options *(continued)*

Field	Description
Logical Characteristics	Displays the logical characteristics that you have chosen to discover. For example, logical disk, logical disk to host mapping, replications, logical groups and so on. The displayed command is the exact command line that the agentlet uses to discover the array object type.
Formula to be used for constructing device identifier(s)	Displays the formulas for enclosure and logical disks identifiers.
Configuration Parameters	Displays the configurable parameters that are required to configure a new device. For example, displays name, its corresponding CLI option, default value, parameter type, and so on.

See "Creating a Storage Insight plug-in" on page 183.

See "Editing a Storage Insight plug-in" on page 192.

Editing a Storage Insight plug-in

The edit option lets you make changes to the existing Storage Insight plug-in. You need to first upload the existing plug-in file (`.sfa` file) to the host. The Storage Insight SDK Add-on extracts the required information from the file for editing. You can edit the information, and the create the next version of the Storage Insight plug-in.

To edit the Storage Insight plug-in

1. Log on to the host, and launch the Storage Insight SDK Add-on console by typing the following URL in the Web browser:

   https://<SDK-hostname>:5634/admin/htdocs/sisdk.html

2. Select the **Edit SI Plug-in** option, and click **Next**.

   See "Create, edit, and test Storage Insight plug-in panel options" on page 183.

3. Upload the Storage Insight plug-in that you want to edit, and click **Edit**.

   See "Upload Storage Insight Plug-in panel options" on page 193.

4 Edit the plug-in related details in the **Plug-in Vitals** panel, and click **Next**.

See "Storage Insight Plug-in Vitals panel options" on page 184.

5 Edit the configuration parameters in the **Configuration Parameters** panel, and click **Next**.

See "Configuration Parameters panel options" on page 185.

6 Edit the physical, and logical capabilities that you want to discover in the **Enclosure Capabilities** panel, and click **Next**.

See "Enclosure Capabilities panel options" on page 186.

7 Edit the device identifier-related information in the **Device Identifier** panel, and click **Next**.

See "Device Identifier panel options" on page 190.

8 Review the information that you have edited on the **Confirmation** panel, and click **Create Add-on**. After the plug-in is created successfully, click **Download Add On** to download it.

See "Confirmation panel options" on page 191.

See "Testing a Storage Insight plug-in" on page 193.

See "About creating Storage Insight plug-in" on page 155.

Upload Storage Insight Plug-in panel options

Use this wizard panel to upload an existing Storage Insight plug-in that you want to edit. After the successful upload, the Storage Insight SDK Add-on extracts the relevant information from the plug-in. You can edit this information, and create your plug-in. Click **Browse** to navigate to the Storage Insight plug-in file, and then click **Upload**.

See "Editing a Storage Insight plug-in" on page 192.

Testing a Storage Insight plug-in

Before you deploy the Storage Insight plug-in, you can ensure that it functions correctly.

Note: The data discovered during the test is not reported to Veritas InfoScale Operations Manager Management Server, and it is not available on the Veritas InfoScale Operations Manager console.

You can test the plug-in based on the following data source options:

- Test with live data: If you choose this option, it is expected that the vendor CLI, or tool used by the discovery script are installed on the SDK host itself. The Storage Insight plug-in runs the discovery script on the SDK host to get the array data.
- Test with collected data: This option is useful when you do not have the vendor CLI or tools installed on the SDK host, or you do not have the connectivity with the array to be discovered. You can run the discovery script on another host that has the vendor CLI installed, and has connectivity to the array. Copy the output of the discovery commands to `.txt` files. The `txt` files can be added to the `tar.gz` file, and the `tar.gz` file can be uploaded as the test data source. Instead of calling the discovery script directly, the agentlet uses this file to discover the array.

Note: There should be a different txt file for each command output. The name of the txt file should be `Command.txt`. For example, for `--list pdevs` command, the output should be stored in `pdevs.txt` file.

To test the Storage Insight plug-in

1. Log on to the host, and launch the Storage Insight SDK Add-on console by typing the following URL in the Web browser:

   https://<SDK-hostname>:5634/admin/htdocs/sisdk.html

2. Select the **Test SI Plug-in** option, and click **Next**.

   See "Create, edit, and test Storage Insight plug-in panel options" on page 183.

3. Select the plug-in to test, and click **Upload**.

   See "Upload Storage Insight Plug-in panel options" on page 195.

4. Click **Test**.
5. Select the data source for the test. You can select live test, or test with the collected data.
   - Test with live data: For the live data option, enter run-time configuration parameters. For example, user name and password.
   - Test with collected data: For the collected data test, upload the required file. The current version of add-on supports only tar.gz files.
6. Click **Run Test**.

See "Creating a Storage Insight plug-in" on page 183.

See "Editing a Storage Insight plug-in" on page 192.

Upload Storage Insight Plug-in panel options

Use this wizard panel to upload an existing Storage Insight plug-in that you want to test. Click **Browse** to navigate to the Storage Insight plug-in file, and then click **Upload.**

See “Testing a Storage Insight plug-in” on page 193.

Section 7

Storage Provisioning and Enclosure Migration Add-on 7.3

- Chapter 15. Provisioning storage
- Chapter 16. Migrating volumes

Chapter 15

Provisioning storage

This chapter includes the following topics:

- About storage provisioning
- About creating a storage template
- Creating a storage template using VxFS file systems
- Creating a storage template using NTFS file systems
- Creating a storage template using volumes
- Updating a storage template
- Provisioning storage
- Uploading storage templates
- Downloading storage templates
- Deleting storage templates
- Locking storage templates
- Unlocking storage templates

About storage provisioning

Storage provisioning is the process of assigning storage, usually in the form of server disk drive space to optimize the performance of a storage area network (SAN). The storage template is an efficient way to quickly create storage based on the configurations that are mentioned in the template. This mechanism lets you specify storage attributes and conditions, and the template can be used by the storage operators to allocate storage. It ensures that the allocated storage contains

attributes mentioned in the template, and it adheres to the conditions that are specified in the template.

Veritas InfoScale Operations Manager storage provisioning lets you create storage templates from existing virtual storage devices such as volumes and file systems. It extracts all the properties of the virtual storage device, and recreates the entire stack. The supported file systems are Veritas File System (VxFS) and NTFS.

With Veritas InfoScale Operations Manager storage provisioning, you determine the virtual device characteristics of the storage object using the file system, volume manager, and operating system utilities and can create a storage template based on this information. A storage template lets you migrate existing virtual devices to new ones with different properties. For example, you can migrate volumes from thick LUNs to thin LUNs.

You can use the Management Server console to create storage templates and provision storage.

You can also use the Veritas InfoScale Operations Manager Web services API version 6.1 or later to provision storage using a storage template that you previously created.

For information on using the Veritas InfoScale Operations Manager Web services API, refer to the *Veritas Operations Manager Management Server User Guide*.

For information on the Veritas InfoScale Operations Manager versions that the add-on is compatible with, refer to the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

See “About creating a storage template” on page 198.

About creating a storage template

You can create a storage template using one of the following methods:

- Modeling it on an existing VxFS file system
  See “Creating a storage template using VxFS file systems” on page 199.
- Modeling it on an existing NTFS file system
  See “Creating a storage template using NTFS file systems” on page 205.
- Modeling it on an existing Storage Foundation volume
  See “Creating a storage template using volumes” on page 206.

Creating a storage template using VxFS file systems

In the Management Server console, you can create a new storage template using VxFS file systems. However, note that the file systems with RAID 5 volumes are not supported.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create a storage template using VxFS file systems

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
2. In the **Actions** pane, click **Create Template**.
3. In the **Create Storage Template** wizard, choose **Using a NTFS or VxFS file system**, and click **Next**.
4. Select a host, a disk group, and a VxFS file system, and click **Next**.
5. In the **Summary** tab, enter a unique storage template name, and optionally an operating system and a description.
6. Select the file system object from the left navigation tree and on the **FS Summary** tab, enter a unique name for the file system template, and optionally a description. To view and update the file system template parameters, click the **Parameters** and **Advanced** tabs.

   See “Create Storage Template – File system parameters panel options” on page 200.

   See “Create Storage Template – File system advanced panel options” on page 200.
7. Select the volume object from the left navigation tree. Click the **Summary**, **Parameters**, and **Advanced** tabs to view or update the volume name, description and parameters.

   See “Create Storage Template – Volume parameters panel options” on page 203.

   See “Create Storage Template – Volume advanced panel options” on page 204.
8. Click **Finish**. The results are displayed. Click **Close**. You can now use the template to provision storage.

   See “Provisioning storage” on page 208.

Create Storage Template – File system parameters panel options

Use this wizard panel to view or update the storage template file system parameters fields when you create or update the storage template.

Table 15-1 Create Template - File system parameters

Field	Description
Size	File system size.
Ask	Lets you change the value while provisioning the storage.
Create file system	Select this check box to create a file system while provisioning the storage.
Mount point prefix	Enter the file system mount point prefix.
Ask	Lets you set the mount point prefix while provisioning the storage.
Block size	File system block size in bytes.
Windows NTFS Parameters	
Volume Label	Enter the volume label for Windows NTFS.
Quick Format	Specify if Quick Format is allowed for the volume.
Compression	Specify if data compression is allowed for the volume.
Allocation Unit Size	Select the required allocation unit size for the volume.

See "Creating a storage template using VxFS file systems" on page 199.

See "Creating a storage template using volumes" on page 206.

See "Updating a storage template" on page 207.

Create Storage Template – File system advanced panel options

Use this wizard panel to view or update the storage template file system advanced fields when you create or update the storage template. These parameters apply to running the `mount` and `vxtunefs` commands.

Note: When you create a storage template based on Veritas File System (VxFS), you can set all the parameters. However, some of these parameters are not supported on certain platforms.

Table 15-2 Create Storage Template - Performance Tunables Tab

Field	Description
Intent log size (bytes)	Enter the size of the intent log in the multiples of size byte.
Intent log behavior (mount option)	Select an option to control the timing of flushing the VxFS intent log and other metadata to disk. This affects when operations are guaranteed persistent after a system failure.
I/O feature (mount option)	Select the required mount I/O feature.
Read ahead	Select the required option. This implements an algorithm that allows read ahead to detect more elaborate patterns (such as increasing or decreasing read offsets, or multithreaded file accesses) in addition to simple sequential reads.
Allocation unit size (bytes)	Enter the size of the allocation unit.
Initial extent size (blocks)	Enter a value for the initial extent.
Maximum extent size (bytes)	Enter the maximum size of an extent.
HSM write preallocation	Select the required option. For a file managed by a Hierarchical Storage Management (HSM) application, this option preallocates disk blocks before data is migrated back into the file system.
Maximum disk queue (bytes)	Enter a value for the maximum disk queue that is generated by a single file.
Write throttle	Enter a value to lower the number of dirty pages per file that the file system generates before writing them to disk.

Table 15-3 File Change Log Tunables Tab

Field	Description
FCL keep time	Enter a value for the VxFS File Change Log (FCL) to keep records in the log.
FCL maximum allocated space	Enter the maximum amount of space that can be allocated to the VxFS File Change Log.
FCL open interval	Enter the time interval in seconds within which subsequent opens of a file do not produce an additional FCL record.
FCL write interval	Enter the time, in seconds, that must elapse before the VxFS File Change Log records a data overwrite, data extending write, or data truncate for a file.

Table 15-4 Other Options Tab

Field	Description
Auto mount	Select an option. This lets the file system to be mounted explicitly.
Preferred read request size	Enter the preferred read request size.
Number of parallel read requests	Enter the number of parallel read requests of size read_pref_io that can be outstanding at one time.
Preferred write request size	Enter the preferred write request size.
Number of parallel write requests	Enter the number of parallel write requests of size write_pref_io that can be outstanding at one time.
Maximum inode aging count	Enter the maximum number of inodes to place on an inode aging list.
Minimum inode size	Enter the minimum size to qualify a deleted inode for inode aging.
Large files enabled	Select an option to make the file system is large files capable.

See “Creating a storage template using VxFS file systems” on page 199.

See “Creating a storage template using volumes” on page 206.

See "Updating a storage template" on page 207.

Create Storage Template – Volume parameters panel options

Use this wizard panel to define or update the volume parameters in a storage template when you create or update the storage template.

Table 15-5 Create Storage Template - Volume parameters options

Field	Description
Layout	Volume layout, which can be **Concat**, **Concat Mirror**, **Mirror Concat**, **Stripe**, **Stripe Mirror**, **Mirror Stripe.**
# Columns	The number of columns (for stripe volumes).
# Mirrors	The number of mirrors (for mirrored volumes).
Name prefix	The volume name prefix.
Minimum number of paths	The minimum number of DMP paths for the disks on which the volume is created.
Fast resync	Turn on/off the Storage Foundation fast resync option.

In the **Rules for selecting LUNs** area, you can view, add, and edit rules for selecting the LUNs. If you have specified mirrored layout options, tabs are displayed corresponding to the number of mirrors that you have specified.

For each mirror, you can define the rules that Veritas InfoScale Operations Manager can use to select LUNs for hosting the mirror.

- To add a rule, click the plus (**+**) icon at the bottom of the rules list.
- To remove a rule, click the minus (**-**) icon.
- To edit an existing rule, double-click it. The **Select LUN Characteristics** dialog box is displayed.
  See "Select LUN Characteristics options" on page 204.

To create mirror plex on multiple enclosures, select the **Mirror across enclosures** check box.

To create stripe volumes on multiple enclosures, select the **Stripe across enclosures** checkbox.

See "Creating a storage template using VxFS file systems" on page 199.

See "Creating a storage template using volumes" on page 206.

See “Updating a storage template” on page 207.

Select LUN Characteristics options

Use this dialog box to view or edit required LUN characteristics for a storage template.

Table 15-6 Select LUN Characteristics options

Field	Description
Enclosure VDID	Displays the vdid of the enclosure.
Vendor	Select an enclosure vendor.
Product	Select an enclosure product name.
LUN Type	Select the type required, for example, **Thin**, **Thick**, or **Any** (the default).
Tier	Select a predefined LUN classification. LUN classifications are defined by the storage administrator on the **Storage** perspective. See “About LUN classification” on page 106.
Media	Select the type of the media, for example, SSD (solid-state drive) or HDD (hard disk drive), or **Any** (the default).
Replicated	Select whether to use replicated LUNs.

See “Creating a storage template using VxFS file systems” on page 199.

See “Creating a storage template using volumes” on page 206.

See “Updating a storage template” on page 207.

Create Storage Template – Volume advanced panel options

Use this wizard panel to view or update the storage template volume advanced fields upon creation or update of the storage template.

Table 15-7 Create Storage Template - Volume advanced options

Field	Description
Mode	Enter a value to set the read or write access on the volume.

Table 15-7 Create Storage Template - Volume advanced options *(continued)*

Field	Description
Exclusive	Select **On** to specify that only one node in the cluster can open an existing volume at a time. Select **Off** to specify that more than one node in a cluster can open a volume simultaneously

See "Creating a storage template using VxFS file systems" on page 199.

See "Creating a storage template using volumes" on page 206.

See "Updating a storage template" on page 207.

Creating a storage template using NTFS file systems

In the Management Server console, you can create a new storage template using NTFS file systems.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create a storage template using NTFS file systems

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
2. In the **Actions** pane, click **Create Template**.
3. In the **Create Storage Template** wizard, choose **Using a NTFS or VxFS file system**, and click **Next**.
4. Select a host, a disk group, and a NTFS file system, and click **Next**. .
5. In the **Summary** tab, enter a unique storage template name, and optionally an operating system and a description..
6. Select the file system object from the left navigation tree. Click the **Summary** and **Parameters** tabs to view or update the file system name, description, and parameters.

   See "Create Storage Template – File system parameters panel options" on page 200.

7 Select the volume object from the left navigation tree. Click the **Summary** and **Parameters** tabs to view or update the volume name, description, and parameters.

See "Create Storage Template – Volume parameters panel options" on page 203.

8 Click **Finish**. The results are displayed. Click **Close**. You can now use the template to provision storage.

See "Provisioning storage" on page 208.

Creating a storage template using volumes

In the Management Server console, you can create a storage template using volumes. Storage templates let you store the volume configuration information in the form of a template so that you can use them on other volumes.

Storage Provisioning and Enclosure Migration Add-on supports only Storage Foundation volumes. RAID5 volumes are not supported.

Note: Only hosts with VxFS volumes are displayed on the host table. Depending on which host is selected, the disk group choices change. Only disk groups with Storage Foundation volumes are displayed on the disk group table. Depending on which host and disk group are selected, the volume choices change.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create a storage template using volumes

1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.

2 In the **Actions** pane, click **Create Template**.

3 In the **Create Storage Template** wizard, choose **Using a volume**. Click **Next**

4 Select a host, a disk group, and a volume, and click **Next**.

5 In the **Summary** tab, enter a unique storage template name and optionally an operating system and a description. .

6 Select the file system object from the left navigation tree. Click the **Summary**, **Parameters**, and **Advanced** tabs to view or update the file system name, description and parameters.

See "Create Storage Template – File system parameters panel options" on page 200.

See "Create Storage Template – File system advanced panel options" on page 200.

7 Select the volume object from the left navigation tree. Click the **Summary**, **Parameters**, and **Advanced** tabs to view or update the volume name, description and parameters.

See "Create Storage Template – Volume parameters panel options" on page 203.

See "Create Storage Template – Volume advanced panel options" on page 204.

8 Click **Finish**. The results are displayed. Click **Close**.

Updating a storage template

In the Management Server console, you can update or view a storage template.

To perform an update task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To update a storage template

1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.

2 In the table, right-click the storage template and select **Update Template**.

3 In **Update Template** wizard, optionally modify the template information in the **Summary** tab. The storage template name is required and it must be unique.

4 Expand the storage template in the left navigation to view its associated volume and file system.

5 Select the file system from the left navigation tree. You can view or update the file system-related information in the following tabs:

- **FS Summary** tab
  You can update the name and description.
- **Parameters** tab
  See "Create Storage Template – File system parameters panel options" on page 200.

- **Advanced** tab: This tab is displayed only if you selected the **Create File System** check box in the **Parameters** tab.
  See "Create Storage Template – File system advanced panel options" on page 200.

6 Select the volume from the left navigation tree. You can view or update the volume-related information in the following tabs:

- **Volume Summary** tab
- **Parameters** tab
  See "Create Storage Template – Volume parameters panel options" on page 203.
- **Advanced** tab
  See "Create Storage Template – Volume advanced panel options" on page 204.

7 To save the updated storage template information, click **Finish**. The results are displayed. Click **Close**.

Provisioning storage

In the Management Server console, you can provision storage using a storage template.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To provision storage

1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.

2 In the **Actions** pane, click **Provision Storage**.

3 Select a host and a disk group that the storage should be provisioned on. (When you select a host, the associated disk groups are displayed for your selection.) Click **Next**.

4 Select a storage template.

If **Ask** is selected in the storage template for the size or for the mount point, you can edit the corresponding fields.

5 Click **Finish** to provision storage. The task information is displayed so that you can track the task in the Tasks pane. Click **Close.**

Uploading storage templates

In the Management Server console, you can upload storage templates from your computer to the Veritas InfoScale Operations Manager repository. You can upload one template at a time.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To upload a storage template

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
2. In the **Actions** pane, click **Upload Template.**
3. In **Upload Templates** panel, click **Load template file** to navigate to the template that you want to upload and click **Open**.
4. Click **Upload** to upload the template to the Veritas InfoScale Operations Manager repository.

See "Downloading storage templates" on page 209.

Downloading storage templates

In the Management Server console, you can download storage templates. You can select multiple storage templates and download them simultaneously. When you download the storage templates, they are saved on your computer in a single file. The downloaded file has a `.json` (JavaScript Object Notation) extension. For example, if you download two storage templates, the application creates a combined downloaded file (`2tmpls.json`) that you need to save on your computer.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To download one or more storage templates

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
2. In the **Actions** pane, click **Download Template.**
3. In the **Download Templates** panel, select the templates to download and click **Download**.

See "Uploading storage templates" on page 209.

Deleting storage templates

In the Management Server console, you can delete storage templates that you created.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To delete one or more storage templates

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
2. Right-click the storage template and select **Delete Template**.
3. In the **Delete Storage Template** panel, confirm the operation. You can delete only those storage templates that are editable, and that you have created.

Locking storage templates

In the Management Server console, you can lock storage templates that you have created.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To lock one or more storage templates

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
2. Right-click the storage template and select **Lock Storage Template.**
3. In the **Lock Storage Template** panel, click **Yes** to lock the storage template. You can lock only those storage templates that you have created.

See “Unlocking storage templates” on page 210.

Unlocking storage templates

In the Management Server console, you can unlock storage templates that you have created.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To unlock one or more storage templates

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
2. Right-click the storage template and select **Unlock Storage Template**.
3. In the **Unlock Storage Template** panel, click **Yes** to unlock the storage template. You can unlock only those storage templates that you have created.

See “Locking storage templates ” on page 210.

Chapter 16

Migrating volumes

This chapter includes the following topics:

- About volume migration
- About the Impact Analysis report for volume migration
- Migrating volumes by host
- Select LUN Characteristics options
- Change layout options
- Migrating volumes by enclosure
- Migrating volumes by disk group
- Pausing or resuming a volume migration
- Deleting a scheduled volume migration

About volume migration

Enterprises frequently need to move (migrate) storage volumes to new disk enclosures. Volume migration can be time-consuming and complex. Storage Provisioning and Enclosure Migration Add-on simplifies the complexities so that you can easily migrate volumes between disk enclosures.

Using the Management Server console, you can migrate volumes in the following ways.

See “Migrating volumes by host” on page 214.

See “Migrating volumes by enclosure” on page 217.

See “Migrating volumes by disk group” on page 219.

Note: Volume migration is supported for managed hosts that run Veritas InfoScale Operations Manager (VRTSsfmh package) 5.x and later.

Before migrating volumes from one enclosure to the target disks in another enclosure, consider the following factors:

- The volumes to be migrated
  Storage Provisioning and Enclosure Migration Add-on migrates only the Storage Foundation-managed volumes.
  A storage array might contain volumes that are managed by a different volume manager than Storage Foundation. It might also contain disks directly used by the hosts. The add-on cannot migrate such volumes and will ignore them.
- The target enclosure
  Ensure that the target enclosure is connected to the host to which the source volumes are connected. Both the disks that contain the volumes to be migrated and the target disks must be accessible from the same host.
- The target disks where you want to migrate the volumes
  The size and layout of selected volumes determine the number of target disks you need. Ensure that the size and number of the target disks are adequate to migrate the selected volumes. Ensure that you add the disks to the disk group. If you use the disk group migration wizard, you can add the disks using the migration wizard.

Volume migration stops if the host shuts down while migrating the selected volumes to the specified target disks. However, it does not result in loss of data. You can start a new volume migration case to migrate the remaining volumes.

Veritas InfoScale Operations Manager lets the storage administrator define LUN classifications. LUN classifications can specify the LUN characteristics, such as RAID levels and replication, that match the SLAs for the applications that are using the volumes. If LUN classifications are defined, you can use them to determine the target LUNs when you set up storage migration.

See “About LUN classification” on page 106.

About the Impact Analysis report for volume migration

When you migrate volumes using Storage Provisioning and Enclosure Migration Add-on, the migration wizard can generate an Impact Analysis report, in addition to the migration summary displayed in the wizard. The impact analysis shows additional information about the objects that are associated with the volumes being

migrated. You can analyze the potential outcome of the volume migration and if necessary change the options you have selected before executing the volume migration.

You can view information such as the following.

- The details about the databases that may be affected by the volume migration.
- The details about Business Applications that may be affected by the volume migration.

See "About volume migration" on page 212.

Migrating volumes by host

You can migrate volumes that reside on an enclosure to one or more other enclosures. This topic covers migration by host. You can also migrate by enclosure or disk group. Disk group migration is available only on the UNIX/Linux hosts.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Before you migrate volumes, review the planning information in the following:

See "About volume migration" on page 212.

In addition, verify the following:

- If volumes have been set up for creating snapshots, unprepare those volumes.
- Ensure that disks are provisioned on the new enclosures and are added to the disk group.
- Ensure that the enclosures are added to Veritas InfoScale Operations Manager.

To migrate volumes by host

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Migration**.
2. In the **Actions** pane, click **Migrate Volumes By Host**.
3. In the **Migrate Volumes By Host** wizard, specify a name and description for the migration task so that you can track the task status. Click **Next**.
4. Select the host and then select the source enclosure. Click **Next**.

5 Select one or more enclosures as the targets to which you want to migrate the volumes.

Under **Select LUN Characteristics**, specify the LUN requirements.

See "Select LUN Characteristics options" on page 216.

Click **Next.**

6 Select the volumes to migrate. Verify the source and target layout. If you want to change the target layout, double-click the volume and select the new layout on the **Change Layout** window.

See "Change layout options" on page 216.

You can also choose from the following options:

- Keep the same number of paths to the volume (the default) or set a minimum number of paths.
- Optionally, select the check box to delete the original volumes after successfully migrating the volumes to the target enclosure.

Click **Next.**

7 Choose one of the following:

- Schedule the time of the migration.
- To run immediately, select **Run now**.

Warning: Review the information under **Migration Summary** before starting the migration. For more information, click **Impact Analysis**. A separate window opens to display impact analysis information regarding affected objects.

---

8 When done reviewing the **Migration Summary** and **Impact Analysis**, click **Previous** if you want to go back and change anything. Otherwise, click **Finish** and close the confirmation window.

Scheduled migrations are listed on the **Migration Status (By Schedule)** tab. You can check on the status of a specific volume on the **Migration Status (By Volume)** tab.

See "Migrating volumes by enclosure" on page 217.

See "Migrating volumes by disk group" on page 219.

Select LUN Characteristics options

When migrating volumes, you can specify required LUN characteristics for the target enclosure.

Table 16-1 Select LUN characteristics options

Field	Description
Enclosure VDID	Displays the vdid of the enclosure.
Vendor	Select an enclosure vendor. The default is the selected enclosure vendor.
Product	Select an enclosure product name. The default is the selected enclosure product name.
LUN Type	Select the type required, for example, **Thin**, **Thick**, or **Any** (the default).
Classification	Optionally, select a predefined LUN classification. LUN classifications are defined by the storage administrator on the **Storage** perspective. See "About LUN classification" on page 106.
Media	Select the type of the media, for example, SSD (solid-state drive) or HDD (hard disk drive), or **Any** (the default).
Raid Type	Select the RAID type, for example, RAID0.
Replicated	Select whether to use replicated LUNs.

See "Migrating volumes by host" on page 214.

See "Migrating volumes by enclosure" on page 217.

Change layout options

Use this panel to change the volume layout of the target volume when selecting volumes for migration.

Table 16-2 Change volume layout options

Field	Description
Target Layout	Select one of the following: ■ **concat** - A layout style that is characterized by the subdisks that are arranged sequentially and contiguously. ■ **concat-mirror** - A combination of concat and mirror, where mirror represents extra copies of the data for data protection. ■ **mirror -concat** - A combination of mirror and concat. ■ **striped** - A layout technique that spreads data across several physical disks using stripes. The data is allocated alternately to the stripes within the subdisks of each plex. ■ **stripe-mirror** - A combination of stripe and mirror.
No. of mirrors	Select the number of mirrors, for mirrored layouts. The default is 2.
Stripe width	For striped layouts, select the stripe width. The default is 128.
No. of columns	For striped layouts, select the number of columns. The default is 2.

See "Migrating volumes by host" on page 214.

See "Migrating volumes by enclosure" on page 217.

See "Migrating volumes by disk group" on page 219.

Migrating volumes by enclosure

You can migrate volumes that reside on an enclosure to one or more other enclosures. This topic covers migration by enclosure. You can also migrate volumes by host and by disk group. Disk group migration is available only on the UNIX/Linux hosts.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Before you migrate volumes, review the planning information in the following:

See "About volume migration" on page 212.

In addition, verify the following:

- If volumes have been set up for creating snapshots, unprepare those volumes.
- Ensure that disks are provisioned on the new enclosures.
- Ensure that the enclosures are added to Veritas InfoScale Operations Manager.

To migrate volumes by enclosure

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Migration**.
2. In the **Actions** pane, click **Migrate Volumes By Enclosure**.
3. In the **Migrate Volumes By Enclosure** wizard, specify a name and description for the migration task so that you can track the task status. Click **Next**.
4. Select the source enclosure and click **Next**.
5. Select one or more enclosures as the targets to which you want to migrate the volumes.

   Under **Select LUN Characteristics**, specify the LUN requirements and click **Next**.

   See "Select LUN Characteristics options" on page 216.
6. Select a host and click **Next**.
7. Select the volumes to migrate. Verify the source and target layout. If you want to change the target layout, double-click the volume and select the new layout.

   See "Change layout options" on page 216.

   You can also choose from the following options:

   - Keep the same number of paths to the volume (the default) or set a minimum number of paths.
   - Optionally, select the check box to delete the original volumes after successfully migrating the volumes to the target enclosure.

   Click **Next.**
8. Choose one of the following:

   - Schedule the time of the migration.
   - To run immediately, select **Run now**.

Warning: Review the information under **Migration Summary** before starting the migration. For more information, click **Impact Analysis**. A separate window opens to display impact analysis information regarding affected objects.

9 When done reviewing the **Migration Summary** and **Impact Analysis**, click **Previous** if you want to go back and change anything. Otherwise, click **Finish** and close the confirmation window.

Scheduled migrations are listed on the **Migration Status (By Schedule)** tab. You can check on the status of a specific volume on the **Migration Status (By Volume)** tab.

See "Migrating volumes by host" on page 214.

See "Migrating volumes by disk group" on page 219.

Migrating volumes by disk group

You can migrate a disk group that resides on an enclosure to one or more other enclosures. While setting up the migration you can optionally add free, uninitialized disks from the target enclosures to the disk group. This disk group operation is available only on the UNIX/Linux hosts and requires Storage Provisioning and Enclosure Migration Add-on.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Before you migrate volumes by disk group, review the planning information in the following:

See "About volume migration" on page 212.

In addition, verify the following:

- If volumes have been set up for creating snapshots, unprepare those volumes.
- Ensure that disks are provisioned on the new enclosures. You can use the migration wizard to add disks to the disk group.
- Ensure that the enclosures are added to Veritas InfoScale Operations Manager.

To migrate volumes by disk group

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Expand the host and expand **Disk Groups** to locate the disk group.

4 Right-click the disk group and click **Migrate Storage**.

5 In the **Migrate Volumes by Disk Group** wizard, specify a name and description for the migration task so that you can track the task status. Click **Next**.

6 Select the source and target enclosure. Optionally, to add disks to the disk group, select the check box **Add free/uninitialized disks to the disk group**. Click **Next**.

7 If you selected the option to add disks, select one or more disks to add to the disk group from the list of available free (initialized) or uninitialized disks. Click **Next**.

8 Select the volumes to migrate. Verify the source and target layout. If you want to change the target layout, double-click the volume and select the new layout on the **Change Layout** window.

See "Change layout options" on page 216.

You can also choose from the following options:

- Keep the same number of paths to disks for selected volumes (the default) or set a minimum number of paths.
- Optionally, select the check box to delete the original volumes after successfully moving the volumes to the target enclosure.

Click **Next.**

9 Choose one of the following:

- Schedule the time of the migration.
- To run immediately, select **Run now**.

Warning: Review the information under **Migration Summary** before starting the migration. For more information, click **Impact Analysis**. A separate window opens to display impact analysis information regarding affected objects.

Click **Next**.

10 When done reviewing the **Migration Summary** and **Impact Analysis**, click **Previous** if you want to go back and change anything. Otherwise, click **Finish** and close the confirmation window.

Scheduled migrations are listed in the **Solutions** area for the **Storage Migration** solution on the **Migration Status (By Schedule)** tab. You can check on the status of a specific volume on the **Migration Status (By Volume)** tab.

See “Migrating volumes by host” on page 214.

See “Migrating volumes by enclosure” on page 217.

Pausing or resuming a volume migration

If a volume migration task is not complete, you can pause the migration and later resume it. The Pause/Resume operation is only available on a UNIX/Linux managed host with Veritas InfoScale Operations Manager version 6.0 or later.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To pause a volume migration

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Migration**.
2. In the **Migration Status (by Schedule)** tab, right-click the scheduled migration task and click **Pause**.
3. In the **Pause Migration** wizard, select the volumes to pause and click **Finish**.

To resume a volume migration

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Migration**.
2. In the **Migration Status (by Schedule)** tab, right-click the paused migration task and click **Resume**.
3. In the **Resume Migration** wizard, select the volumes to resume and click **Finish**.

Deleting a scheduled volume migration

If a scheduled volume migration is not in progress, you can delete it.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To delete a scheduled volume migration

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Migration**.
2. In the **Migration Status (by Schedule)** tab, right-click the scheduled migration task and click **Delete Schedule**.

Section 8

Veritas HA Plug-in for VMware vSphere Web Client

- Chapter 17. Introduction to Veritas HA Plug-in for vSphere Web Client
- Chapter 18. Installation and uninstallation of Veritas HA Plug-in for vSphere Web Client
- Chapter 19. Configurations for Veritas HA Plug-in for vSphere Web Client

Chapter 17

Introduction to Veritas HA Plug-in for vSphere Web Client

This chapter includes the following topics:

- About Veritas HA Plug-in for vSphere Web Client
- Limitations of Veritas HA Plug-in for vSphere Web Client

About Veritas HA Plug-in for vSphere Web Client

The Veritas HA Plug-in for vSphere Web Client lets you administer Veritas application availability products and features directly from VMware vSphere Web Client. It provides a VMware vCenter plug-in to integrate the high availability products from Veritas with the VMware vSphere Web Client. With this plug-in you can use the client to monitor applications running inside the virtual machines that are under a VMware vCenter Server control.

Veritas HA Plug-in for vSphere Web Client supports the integration of the following Veritas products with vSphere Web Client:

- Veritas AppProtect: Just In Time availability of applications during planned and unplanned maintenance of virtual machines.
  For details, see the *Veritas InfoScale Solutions Guide*.
- Cluster Server: Application high availability monitoring at a virtual machine or virtual machine cluster level from the High Availability tab. Application high availability monitoring at an ESX cluster level, or data center level, from the High Availability Dashboard. The add-on also supports monitoring for the applications that are configured for disaster recovery with VMware Site Recovery Manager.

For details, see the Cluster Server documentation.

- ApplicationHA: Application availability monitoring at a virtual machine level from the High Availability tab. Application high availability monitoring at an ESX cluster level, or data center level, from the High Availability Dashboard.
  For details, see the ApplicationHA documentation.

New users as well as users migrating from vSphere Client (desktop) can use Veritas HA Plug-in for vSphere Web Client. Existing users of vSphere Client (desktop) can optionally continue to use both the desktop client and the Web client.

You can also deploy the Veritas HA Plug-in for vSphere Web Client in a high availability configuration of the Management Server.

Refer to the following for more information:

- For a list of supported Veritas InfoScale Operations Manager versions and compatible add-ons, see the compatibility list at:
  https://www.veritas.com/support/en_US/article.000107928
- See "Registering the HA Plug-in with VMware vCenter Server" on page 228.
- See "Deploying the HA Plug-in if the Management Server is configured in a high availability environment" on page 230.
- See "Unregistering the HA Plug-in from VMware vCenter Server" on page 229.

Limitations of Veritas HA Plug-in for vSphere Web Client

This section lists the limitations of Veritas HA Plug-in for vSphere Web Client:

- Ability to install ApplicationHA guest component via vSphere Web Client: You cannot use the vSphere Web Client to install the ApplicationHA guest components on the virtual machine. Use the vSphere Desktop Client to install the ApplicationHA guest components.
- The High Availability home page will not be available on the vSphere Web Client, and the operations related to License Management, single sign-on configuration between sites (for VMware Site Recovery Manager) are not available.
- Integration with Backup Exec is not supported.

See "About Veritas HA Plug-in for vSphere Web Client" on page 223.

Chapter 18

Installation and uninstallation of Veritas HA Plug-in for vSphere Web Client

This chapter includes the following topics:

- Before installing the Veritas HA Plug-in for vSphere Web Client
- Installing Veritas HA Plug-in for vSphere Web Client on the Management Server
- Before uninstalling the Veritas HA Plug-in for vSphere Web Client
- Uninstalling the Veritas HA Plug-in for vSphere Web Client

Before installing the Veritas HA Plug-in for vSphere Web Client

Before installing the Veritas HA Plug-in for vSphere Web Client, consider the following:

- Install Veritas InfoScale Operations Manager Management Server, if it is not already installed in your environment.
  For more information, see the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.
- Install and configure the Veritas InfoScale Operations Manager Control Host Add-on.

For more information, see the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

- Install ApplicationHA guest components or Cluster Server (VCS) on the virtual machines where you want to monitor application availability. In the Veritas InfoScale Operations Manager Virtualization perspective verify that the virtual machines are visible.

Installing Veritas HA Plug-in for vSphere Web Client on the Management Server

You can either install only the Veritas HA Plug-in for vSphere Web Client or you can perform a full installation (install Veritas InfoScale Operations Manager Management Server, add-ons or plug-ins). You can use the Management Server to install the Veritas HA Plug-in for vSphere Web Client.

Note: If you configure Veritas InfoScale Operations Manager Management Server in the high availability (HA) mode, you must install and configure the Veritas HA Plug-in for vSphere Web Client on both the primary and secondary nodes of the Management Server.

You can download and install the plug-in using any of the following ways:

- Download the plug-in from Veritas Services Operations Readiness Tools (SORT) website. Then, in the Management Server console, upload the plug-in to the repository using the **Upload Solutions**, and then proceed with the installation.
- Use the Management Server console (under the **Deployment** section of **Settings**) to download and install the plug-in.

For detailed instructions on downloading, uploading, and installing the Veritas InfoScale Operations Manager add-ons, refer to the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

Note: After you install the plug-in, you must restart the Veritas InfoScale Operations Manager Web Server.

Before uninstalling the Veritas HA Plug-in for vSphere Web Client

Before you uninstall the Veritas HA Plug-in for vSphere Web Client, consider the following:

- Uninstalling the Veritas HA Plug-in for vSphere Web Client does not unregister the HA plug-in from the vCenter servers. You must first unregister the HA plug-in from all the vCenter servers and then uninstall the Veritas HA Plug-in for vSphere Web Client.
  See "Unregistering the HA Plug-in from VMware vCenter Server" on page 229.
- If you have configured Veritas InfoScale Operations Manager Management Server in the high availability (HA) mode, you must uninstall the Veritas HA Plug-in for vSphere Web Client from both the primary and secondary nodes of the Management Server.

Uninstalling the Veritas HA Plug-in for vSphere Web Client

You can uninstall the Veritas HA Plug-in for vSphere Web Client from the Management Server. To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

For the detailed instructions on uninstalling Veritas InfoScale Operations Manager add-ons (plug-ins) and removing them from the repository, refer to the *Veritas InfoScale Operations Manager Management Server Installation and Configuration Guide*.

Chapter 19

Configurations for Veritas HA Plug-in for vSphere Web Client

This chapter includes the following topics:

- Registering the HA Plug-in with VMware vCenter Server
- Unregistering the HA Plug-in from VMware vCenter Server
- Deploying the HA Plug-in if the Management Server is configured in a high availability environment
- Adding managed hosts to the Management Server
- Migrating virtual machines to Veritas InfoScale Operations Manager

Registering the HA Plug-in with VMware vCenter Server

After a successful installation of Veritas HA Plug-in for vSphere Web Client, you can view the **Veritas HA Plug-in for vSphere Web Client** option listed under the **Solutions** section of **Virtualization** perspective in the Management Server console. The **Veritas HA Plug-in for vSphere Web Client** tab displays the VMware vCenter servers that are currently discovered in a Veritas InfoScale Operations Manager environment.

Select the required vCenter Server and use the **Register** option to register the HA Plug-in with the vCenter Server. After the registration, ApplicationHA user interface elements (**Dashboard** and **Application HealthView**) will be visible in vSphere Web

Client. In Veritas InfoScale Operations Manager, only the virtualization administrator has the required privileges to register the plug-in.

To register the HA Plug-in with VMware vCenter Server

1. In the Management Server console, go to the **Virtualization** perspective, expand **Solutions** in the left pane, and click **Veritas HA Plug-in for vSphere Web Client**.
2. Under the **Veritas HA Plug-in for vSphere Web Client** tab, select the required vCenter Server to be registered with the HA Plug-in.
3. Right-click and select **Register**.

   Note that the user details used for the plug-in registration are same as the user details used for the vCenter Server discovery in the Veritas InfoScale Operations Manager Management Server. Also, the user must have the VMware Plug-in registration privilege.
4. Click **Finish** on the confirmation wizard panel to register the HA Plug-in with the VMware vCenter Server.

See "Unregistering the HA Plug-in from VMware vCenter Server" on page 229.

See "About Veritas HA Plug-in for vSphere Web Client" on page 223.

Unregistering the HA Plug-in from VMware vCenter Server

You can use the Management Server console to unregister the Veritas HA Plug-in for vSphere Web Client that is registered with a specific VMware vCenter Server.

To unregister the HA Plug-in from VMware vCenter Server

1. In the Management Server console, go to the **Virtualization** perspective, expand **Solutions** in the left pane, and click **Veritas HA Plug-in for vSphere Web Client**.
2. Under the **Veritas HA Plug-in for vSphere Web Client** tab, select the required VMware vCenter Server.
3. Right-click and select **Unregister**.
4. Click **Finish** on the confirmation wizard panel to unregister the HA Plug-in.

See "Registering the HA Plug-in with VMware vCenter Server" on page 228.

See "About Veritas HA Plug-in for vSphere Web Client" on page 223.

Deploying the HA Plug-in if the Management Server is configured in a high availability environment

This section describes two scenarios for the deployment of Veritas HA Plug-in for vSphere Web Client depending on your existing Veritas InfoScale Operations Manager Management Server implementation.

- When Veritas InfoScale Operations Manager Management Server is configured in a high availability (HA) environment
  In this case, you need to install Veritas HA Plug-in for vSphere Web Client on the primary and secondary nodes. After a successful installation of Veritas HA Plug-in for vSphere Web Client, use the Management Server to register the HA Plug-in with the required VMware vCenter Server.
  See "Registering the HA Plug-in with VMware vCenter Server" on page 228.
- When Veritas InfoScale Operations Manager Management Server is not configured in a high availability (HA) environment
  In this case, if you have already installed the Veritas HA Plug-in for vSphere Web Client and registered it on a VMware vCenter server, and you want to migrate the Management Server to a high availability environment, you need to perform the following steps:
  - Using the Management Server, first un-register the HA Plug-in from the VMware vCenter Server.
    See "Unregistering the HA Plug-in from VMware vCenter Server" on page 229.

    Note: Note that this un-registration process does not affect the periodic discovery of VMware vCenter Servers using Veritas InfoScale Operations Manager.
  - Uninstall Veritas HA Plug-in for vSphere Web Client.
  - Configure the Veritas InfoScale Operations Manager Management Server in a high availability (HA) environment.
  - Install Veritas HA Plug-in for vSphere Web Client on the primary and secondary nodes of the Management Server.
  - Re-register the HA Plug-in with the required VMware vCenter Server.

See "About Veritas HA Plug-in for vSphere Web Client" on page 223.

Adding managed hosts to the Management Server

To monitor applications running inside virtual machines under ApplicationHA or Cluster Server (VCS) control, you must first add the virtual machines (guests) as managed hosts to Management Server.

Before you start adding the managed hosts, ensure that the following conditions are fulfilled:

- There is connectivity between the managed hosts and the Management Server.
- VRTSsfmh 5.0 or later is installed on ApplicationHA guests, and VRTSsfmh 6.1 is installed on VCS guests.
- Veritas InfoScale Operations Manager Management Server 6.1 is installed.

  Note: Veritas InfoScale Operations Manager does not display an error if this condition is not fulfilled.

- The logged-on user has adequate administrative or root privileges to execute ApplicationHA or VCS tasks on the virtual machine.

You can add one virtual machine at a time, or add multiple hosts using a CSV file.

Use any of the following methods to add managed hosts:

- To add one or more managed hosts using the Management Server console
- To add a managed host from vSphere Web Client
- In your setup, if the virtual machines are presently attached to the Symantec HA Console, you can migrate all the monitored virtual machines to Veritas InfoScale Operations Manager.
  See "Migrating virtual machines to Veritas InfoScale Operations Manager" on page 232.

To add one or more managed hosts using the Management Server console

1. In the Home page on the Management Server console, click **Settings**.
2. Do one of the following:
   - Click **Add Hosts** > **Agent**.
   - In the **Settings** tab click **Host**, and then click **Add Hosts** > **Agent**.
3. To manually specify one or more hosts, click **Add Entry** and specify the following host details for each host:
   - Host name: Specify the host name or IP address that you can use to reach the host from the Management Server

- User name: Specify a user name with administrative or root privileges
- Password: Specify the password to log on with the specified user name

To specify multiple hosts using a CSV file, under Advanced, browse to a CSV file to specify multiple hosts, in the following format:

Note: Ensure that the first line is exactly as follows.

```
Host, User, Password
host1, user1, password1
host2, user2, password2
host3, user3, password3
```

4 Click **Import selected file** and click **Finish**.

5 In the Result panel, verify that all the specified managed hosts are added.

To add a managed host from vSphere Web Client

1 Download the gendeploy script from the Management Server.

Note: Only a user with Admin role on the Management Server perspective of Veritas InfoScale Operations Manager can generate the script.

---

2 In the Monitor tab of the vSphere Web Client, navigate to the virtual machine that you want to add as a managed host.

If the virtual machine is not already added as a managed host, the Add Host panel appears.

3 In the Add Host panel, specify the user name and password of the user with administrative privileges on the virtual machine.

In the same panel, also specify the location of the gendeploy script, and then click **Configure**.

Migrating virtual machines to Veritas InfoScale Operations Manager

If you use the VMware vSphere Desktop Client to perform ApplicationHA or Cluster Server (VCS) operations, then to be able to perform same tasks from vSphere Web Client, you must migrate all the required virtual machines from the Symantec HA Console to the Veritas InfoScale Operations Manager Management Server using Veritas HA Plug-in for vSphere Web Client.

The following procedure describes the detailed steps that you must perform.

Note: If you are not an existing Symantec HA Console user, skip this procedure.

To migrate virtual machines to Veritas InfoScale Operations Manager

1. Log on to the Management Server as a user with administrative privileges.
2. Navigate to the `hac_generate_csv.pl` script, and execute the following command:

   For Linux:

   ```
   # installdir/bin/perl installdir/adm/hac_generate_csv.pl --vc-name
   vc-name [--vc-user vc-user --vc-password vc-password]
   --ha-console-server ha-console-server
   ```

   For Windows:

   ```
   # "installdir\bin\perl.exe" "installdir\adm\hac_generate_csv.pl"
   --vc-name vc-name [--vc-user vc-user --vc-password vc-password]
   --ha-console-server ha-console-server
   ```

   Where:

   - "installdir" is the installation directory
   - "vc-name" is the FQDN or IP address of the vCenter Server that monitors the required virtual machines.
   - "vc-user" is a user name that has administrative privileges on the vCenter Server

     This is an optional parameter. You need not specify this value if you have already registered the HA Plug-in for vSphere Web Client with vCenter Server.
   - "vc-password" is the password for the user name specified earlier for the "vc-user" parameter.
   - "ha-console-server" is the FQDN or IP address of the Symantec High Availability Console server to which the virtual machines are already attached.

   The script generates the following four files of comma-separated values (CSV):

   - hostlist.csv

     A list of host names of virtual machines attached to the Symantec High Availability Console server. You must edit this CSV file to add the username and password of a user that has administrative privileges for each virtual machine.

For details, see step 3

- MHalready-added.csv
  A list of host names of virtual machines that are attached to both, the Symantec High Availability Console server and the Veritas InfoScale Operations Manager Management Server. Ignore this list for the migration task.
- unsupported-MHversion.csv
  A list of host names of virtual machines on which a supported version of the VRTSsfmh package does not exist. Upgrade the VRTSsfmh package to version 5.0 or later on ApplicationHA guests, and version 6.1 on VCS guests, and then migrate the virtual machines, starting with step 1.
- unresolv-hostlist.csv
  A list of the host name and IP address of each virtual machine, where one or more of the following conditions occur:
  - The host name is not resolvable
  - The VRTSsfmh package is not present or the Veritas Messaging Service (xprtld) is not running.
  - If a virtual machine has multiple IP addresses associated with it, then multiple entries of such a virtual machine will appear in this list

  Resolve these errors and then migrate the virtual machines.

3 Edit the hostlist.csv to specify the user name and password for each virtual machine. Use only the following format:

```
Host, User, Password
host1, user1, password1
host2, user2, password2
host3, user3, password3
```

4 In the Home page of the Management Server console, click **Settings**.

5 Do one of the following:

Click **Add Hosts** > **Agent**.

OR

In the Settings tab click **Host**, and then click **Add Hosts** > **Agent**.

6 Under **Advanced**, browse to the CSV edited in step 3, and click **Import Selected**, and then click **Finish**.

7 In the result panel, verify that all the required virtual machines are added as managed hosts to Veritas InfoScale Operations Manager Management Server.

Note: If you want to migrate the virtual machines attached to multiple Symantec HA Console servers, you must repeat this procedure for all the Symantec HA Console servers.

Section 9

Application Migration Add-on

- Chapter 20. Introduction to Application Migration Add-on
- Chapter 21. Creating and managing an application migration plan

Chapter 20

Introduction to Application Migration Add-on

This chapter includes the following topics:

- About Application Migration add-on
- Before installing the Application Migration add-on
- Installing the Application Migration Add-on
- Before uninstalling the Application Migration add-on
- Uninstalling the Application Migration add-on

About Application Migration add-on

The Application Migration add-on allows you to migrate applications that are under Cluster Server (VCS) management from one cluster to another. The application migration operation is less complex and can be accomplished with minimal manual intervention. Application migration can be across operating systems, architectures, or virtualization technologies. You can migrate data using:

- VxVM Mirroring
- VVR Replication

You can migrate an application between different:

- Platforms—AIX, Linux, and Solaris
- Environments—Physical-to-physical, physical-to-virtual, virtual-to-virtual, and virtual-to-physical
- InfoScale versions

From Application Migration add-on version 7.3 onwards, you can move data from one cluster to another using Veritas Volume Replicator (VVR) replication. For more information about VVR and to understand the basic concepts, refer to the *Veritas InfoScale Replication Administrator's Guide*.

See "Supported versions and platforms" on page 241.

To migrate an application, you must create an application migration plan using the **Create Migration Plan** wizard. After you create a plan, you must execute the migration plan.

- See "Creating an application migration plan" on page 246.
- See "Executing the application migration plan" on page 270.

The add-on also allows you to:

- Pause and resume the operation for manual verification and correction, if required.
- Integrate custom scripts in the operation as per application requirements.
- Migrate application dependencies.
- Understand source cluster configuration and create target cluster configuration.
- Perform endian changes to the data as per architecture requirements.
- Rehearse the steps before the actual migration operation.

Before installing the Application Migration add-on

To install the Application Migration add-on, you must be familiar with Veritas InfoScale Operations Manager concepts and documentation.

Install Veritas InfoScale Operations Manager Management Server, if it is not already installed in your environment.

For more information, see the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

Installing the Application Migration Add-on

You can either install only the Application Migration add-on or you can perform a full installation (install Veritas InfoScale Operations Manager Management Server and add-ons).

You can use the Management Server to install the Application Migration add-on. Download the add-on from the Veritas Services Operations Readiness Tools (SORT)

website. In the Management Server console, upload the add-on to the repository using the **Upload Solutions** option, and then install.

For the detailed instructions on downloading, uploading, and installing the Veritas InfoScale Operations Manager add-ons, refer to the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

Note: After you install the add-on, you must restart the Veritas InfoScale Operations Manager Web Server.

Before uninstalling the Application Migration add-on

Before you uninstall the Application Migration add-on, ensure that you have deleted all the existing application migration plans. If an application migration plan exists, the add-on uninstallation fails.

Before deleting the application migration plan, you can take a backup of the plan by exporting it. If you reinstall the add-on, you can import this plan and use it.

- See “Exporting application migration plan(s)” on page 273.
- See “Importing application migration plan(s)” on page 274.

Uninstalling the Application Migration add-on

You can uninstall the Application Migration add-on from the Management Server. To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

For the detailed instructions on uninstalling Veritas InfoScale Operations Manager add-ons, refer to the *Veritas InfoScale Operations Manager Management Server Installation and Configuration Guide*.

Chapter 21

Creating and managing an application migration plan

This chapter includes the following topics:

- Supported versions and platforms
- User privileges
- Prerequisites for creating an application migration plan
- VVR Replication: Environment variables used in application migration
- Creating an application migration plan
- Understanding user-defined tasks
- Understanding application migration operations
- Understanding the cleanup operation
- Understanding the tasks executed in each operation
- Validations performed before migration plan execution
- Executing the application migration plan
- Editing an application migration plan
- Deleting application migration plan(s)
- Exporting application migration plan(s)
- Importing application migration plan(s)
- Viewing historical runs

- Viewing properties of an application migration plan
- Application migration logs

Supported versions and platforms

The following platforms are supported:

- AIX
- Linux
- Solaris

Table 21-1 lists the supported product or component and versions:

Table 21-1 Supported version(s)

Product or component	Supported version(s)
VIOM Management Server	7.3
VIOM managed host	■ 7.1 ■ 7.2 ■ 7.3
SFHA SFCFSHA	5.1 SP1 RP4 to 6.2.1
InfoScale Availability InfoScale Storage InfoScale Enterprise	7.0 to 7.3
Cluster Server version is from 5.1 SP1 RP4 to 6.0	Managed host version of the cluster must be 7.3 **Note:** You must uninstall the older VRTSsfmh package and install the newer VRTSsfmh package. If you upgrade the package, the required functionality may not be available.
Cluster Server version is from 6.1 to 7.3	Managed host version of the cluster can be the same version as of the cluster. Veritas recommends upgrading to the 7.3 version.

Licenses required for the managed hosts

Mirroring

- If your product version is lower than 7.0, you must have one of the following licenses:
  Storage Foundation Enterprise Edition with Cluster Server
  Or
  Storage Foundation Enterprise Edition and Cluster Server
- If your product version is 7.0 or higher, you must have one of the following licenses:
  InfoScale Enterprise
  Or
  InfoScale Storage and InfoScale Availability

Replication

- If your product version is lower than 7.0, you must have one of the following licenses:
  Storage Foundation Enterprise Edition with Cluster Server and VVR
  Or
  Storage Foundation Enterprise Edition with VVR + Cluster Server
- If your product version is 7.0 or higher, you must have one of the following licenses:
  InfoScale Enterprise
  Or
  InfoScale Storage and InfoScale Availability

User privileges

Table 21-2 lists the user privilege required in the Availability and Server perspectives for performing the various tasks.

Table 21-2 Tasks and user privileges

Task	Availability perspective: Admin	Availability perspective: Admin	Availability perspective: Guest	Availability perspective: Guest
	Server perspective: Admin	Server perspective: Guest	Server perspective: Admin	Server perspective: Guest
Setup Storage	Yes	No	No	No
Rehearse	Yes	No	No	No
Migrate	Yes	No	No	No
Create	Yes	No	No	No
Edit	Yes	No	No	No
Delete	Yes	Yes	No	No
Import	Yes	Yes	No	No
Export	Yes	Yes	No	No
Historical Runs	Yes	Yes	Yes	Yes
Properties	Yes	Yes	Yes	Yes

Prerequisites for creating an application migration plan

In this document, source cluster refers to the cluster from which the service group managing the application must be migrated and target cluster is the cluster to which the service group managing the application must be migrated to.

Ensure that the following prerequisites are met before creating an application migration plan:

- Add all the nodes of the source and target cluster to Veritas InfoScale Operations Manager (VIOM) as managed hosts.
- Application data must reside on a Veritas Volume Manager (VxVM) disk group.

The add-on migrates application data only on VxVM volumes (DiskGroup or CVMVolDG resources). The add-on does not migrate application data on native disks or other volume managers.

- Source cluster must be in running state.
- VCS and VxVM versions in the target cluster must be same or higher than the version in the source cluster.
- All DiskGroup or CVMVolDG type resources in the service groups being migrated must be online.
- Install and configure VCS cluster on the target managed hosts. The target cluster must also be in running state.
- If Cluster Volume Manager (CVM) is configured in the source cluster, install and configure CVM in the target cluster also.
- In the target cluster configuration file, include type definition files for the resource types that are being migrated.
- On the source cluster, all service groups must be online.
- On the managed hosts of the source and target clusters, the XPRTLD daemon must be always running.
- On the target cluster, install and configure the application that is being migrated according to the application and operating system (OS) requirements.
- Ensure that the application has proper service group dependencies because application migration from source to target cluster happens by migrating the service groups in which the application is configured. Dependencies for the selected service group are automatically included as a part of the migration plan.
- Execute the `vxlist` (/opt/VRTSsfmh/bin/vxlist) utility in the source and target cluster node. The output must not contain any errors.
- On the target cluster, identify the IPs that the application will use.
- VVR Replication—On the source and target cluster nodes, identify IPs to create replication links. Ensure that these IPs are not used.
- Depending on the Data Migration type you select, the migration plan detects any pre-existing linked break-off snapshot or VVR replication configured for each disk group between the nodes in the source and target clusters. If no pre-existing configuration is detected, the migration plan creates it and you must identify the list of disks to be used to create either the mirror disk group or remote site disk group. The disks must be under VxVM control and must have enough free space. The total size of the disks you provide must be more than the used

size of the disk group in the source cluster. The number of disks must be at least same or more considering the volume layout in the source cluster.

- VxVM Mirroring—Disks must be shared across all the nodes in the source and target cluster.
- VVR Replication—Disks must be shared between all nodes in the target cluster only.

VVR Replication: Environment variables used in application migration

In the source and target cluster nodes, configure every disk group that must be migrated with a Storage Replicator Log (SRL) volume. During Rehearse operation, create cache volume in every disk group on the secondary site so that space-optimized snapshots can be created. The application migration plan creates the required SRL volume and cache volume. For creating these volumes, 10% of the used size of the disk group is taken as the default size.

To override the default size of both SRL volume and cache volume, mention the size required for a disk group using specific variables in the `appmig.conf` file. The file is not present by default. On the node where the disk group is online, create the file in the `/opt/VRTSsfmh/etc` directory. In the file, define the variable for each disk group.

For example, for a set of service groups to be migrated, if one disk group is online on one node and another disk group is online on another node, create the file for each node and define the variable in the file where the corresponding disk groups are online.

Veritas recommends keeping one copy of the environment file with all necessary entries and copy the same file across all nodes in the source and target clusters.

Note: Mention the size in megabytes.

Define the following environment variables in the file:

- SRLVOL_SIZE_*DiskGroup*
- CACHEVOL_SIZE_*DiskGroup*

Example

To specify 1 GB size for SRL volume and 2 GB size for cache volume for disk group dg1, define the variables as follows:

- SRLVOL_SIZE_dg1 = 1024M
- CACHEVOL_SIZE_dg1 = 2048M

Scenarios in which the environment file (if created) is read

- During plan creation from the source cluster node the file checks whether:
  - Sufficient space is available in the source disk group for creating SRL volume.
  - Sum of the size of disks selected for creating disk group in the target cluster is greater than the sum of the used space in the source disk group, plus SRL volume and cache volume size calculated by the plan.
- During the Setup Storage operation, the SRL volume size is retrieved from the file in the source cluster node and the size is used to create the SRL volume for the disk group in the source and target cluster nodes.
- During the Rehearse operation, the cache volume size is retrieved from the file in the target cluster node and the size is used to create the cache volume for the disk group in the target cluster node.

Creating an application migration plan

Table 21-3 describes the fields in the Create Migration Plan wizard.

Table 21-3 Create Migration Plan wizard options

Field	Description
Components Selection	
Name	Enter a name for the migration plan. The maximum length of the name should be 128 characters. The following characters are allowed: ■ A to Z ■ a to z ■ 0 to 9 ■ _
Description	Enter a description or any additional information for the migration plan.
Select data migration type	Select one of the following: ■ VxVM Mirroring ■ VVR Replication

Table 21-3 Create Migration Plan wizard options *(continued)*

Field	Description
Select a service group to migrate using this plan	Click **Browse** and select the service group that you plan to migrate using this plan. Veritas recommends that the service group you select for migration must be online in the source cluster
Select a target cluster to migrate to using this plan	Click **Browse** and select the target cluster to which you plan to migrate to using this plan.
Cluster Information & Dependency	
Cluster Information	Displays the source and target cluster details. The following details are displayed: ■ OS version ■ VCS version ■ VxVM version **Note:** VCS and VxVM versions in the target cluster must not be lower than the version in the source cluster.
Use existing service group(s) in the target cluster	Option is enabled and selected if the target cluster has service groups of the same name that are part of the entire dependency tree for migration from source.
Create new service group(s) in the target cluster	Option is enabled and selected if there are no service groups or resource names in the target cluster that matches the name of the service group or the resources present in the entire service group dependency tree selected for migration. In this scenario, the application migration plan creates all the service groups selected as a part of the migration. For Cluster Volume Manager (CVM) migrations, this option is enabled even if the CVM group is present in the target cluster as well as in the service group dependency tree which is to be migrated. Veritas recommends that you must not deselect the CVM service group during migration. If you opt to deselect the CVM service group, dependency with CVM service group is not created in the target cluster and may affect while bringing the service group online with correct dependency.
Dependent Service Groups for Migration	You can deselect any dependent service group(s) that need not be migrated. However, Veritas does not recommend deselecting any service group that is in the dependency tree. This may affect the application when it is brought online in the target cluster.

Table 21-3 Create Migration Plan wizard options *(continued)*

Field	Description
Target Disk Information—Panel appears if you select VxVM Mirroring	
Service Group to be Migrated	Displays the name of the service group which has a CVMVolDg or DiskGroup type resource configured in it.
Disk Group	Displays the name of the disk group. Only disk groups configured with CVMVolDg or DiskGroup resources are displayed. If a disk group is not configured under any resource, that disk group is not displayed.
Action	■ Use Existing—Option is enabled only if there are existing linked break-off snapshots for any source disk groups part of the migration. If snapshot is found for any disk group, the corresponding disk group is shown as **Mirrored**. Select this option if you want to use the mirror disk group pre-configured for migration. ■ Create New—Option allows you to create a new mirror disk group for a disk group. Option is enabled by default if there are no pre-configured mirror disk groups for a disk group. In this scenario, the corresponding disk group is marked as **Not Mirrored**. ■ Ignore—You can ignore a particular disk group resource from being migrated. Veritas does not recommend to ignore any disk group as the data present in these disk groups will not be migrated.
Sub-Action	Select Disks—If you opt to create a new disk group, click **Select Disks** to select the set of disks for creating the mirror disk group. The **Select Disk(s)** window displays only the set of disks in the source cluster that are shared between all the target cluster managed hosts. You can select multiple disks so that the combined space of the selected disks is greater than the source disk group space; this is validated when you click **Next**.
RVG Information—Panel appears if you select VVR Replication	
Service Group to be migrated	Displays the name of the service group which has a CVMVolDg or DiskGroup type resource configured in it.
Disk Group	Displays the name of the disk group. Only disk groups configured with CVMVolDg or DiskGroup resources are displayed. If a disk group is not configured under any resource, that disk group is not displayed.

Table 21-3 Create Migration Plan wizard options *(continued)*

Field	Description
Action	■ Use Existing—Option is auto-selected if the disk group has existing RVGs and for all RVGs secondary sites are added. The secondary site for all RVGs should be one of the nodes of the selected target cluster. If a disk group is marked as **Use Existing**, the workflow does not take any action on these disk groups and you must establish replication for these disk groups to ensure data sync. This is validated during Rehearse and Migrate operations. ■ Create New—If a disk group is marked as **Create New**, the workflow creates the necessary entities on the source and target cluster nodes for replication to be established for that disk group. Option is auto-selected for the disk group when one of the conditions is met: ■ Has no RVG created. ■ Has RVGs created but no secondary added. ■ Has RVGs created and few or all RVGs has secondary sites added. None of the secondary site of any RVG should be a node which is part of the target cluster. ■ Ignore—If a disk group is marked as **Ignore**, the disk group is not included for migration as part of the workflow. Option is auto-selected for the disk group when one of the conditions is met: ■ If RVG does not exist in the disk group and the disk group does not have sufficient space to create SRL volume. ■ Has RVGs created and few or all RVGs have a secondary site added, out of which few RVGs have secondary site as one of the nodes of the selected target cluster and few RVGs have secondary site as a node which is not part of the selected target cluster. Configuration issues mentioned above which caused the disk group to be marked as **Ignore** can be resolved as follows: ■ For the first scenario, either increase the size of the source disk group or define the required SRL volume size in the environment file according to the free space available in the disk group. ■ For the second scenario, maintain consistent configuration across all RVGs in a disk group. After resolving the configuration issues, click **Previous** to go back to the previous page and then click **Next** to revalidate. **Note:** You can also ignore a disk group as part of the migration process, but Veritas does not recommend this.

Table 21-3 Create Migration Plan wizard options *(continued)*

Field	Description
Sub-Action	Select Disks—If a disk group is marked as **Create New**, disk information must be provided. Click **Select Disks** for each disk group to provide the target disks on which the disk group for replication needs to be created. **Note:** Combined space of disks selected should be greater than the sum of space used in the source disk group, calculated SRL volume size, and cache volume size. You must not select the same disk as part of two disk groups.
RLink Details	RLink Details—If a disk group is marked as **Create New**, RLink details must be provided. Click **RLink Details** to provide RLink information for each disk group: ■ NIC for each node of source and target cluster where the IP for RLink needs to be plumbed. ■ IP which is to be used for the disk group on source cluster and target cluster for setting up RLinks for replication. **Note:** The NIC selected for a disk group for all nodes of source and target cluster should have same version of pre-existing IP plumbed on it (either IPv4 or IPv6). IP which is provided for a source cluster and target cluster should be unique for disk groups across service groups. Same IP can be given for disk groups which are part of the same service group.
Network Configuration	
Service Group Name	Displays the name of the service group under which network resources are configured.
Resource Name	Displays the name of IP-related resources.
Network Parameters	The following are the network parameters: ■ Device ■ IP ■ Netmask
Source Details	Displays the network details of the source.

Table 21-3 Create Migration Plan wizard options *(continued)*

Field	Description
Target Details	Displays the network details of the target. You can edit the target network details. The target details must be in the same format as the source. As an example, for IPMultiNIC resource, if NIC field displays `net1,1,net2,2` in source, the target information that you provide must be in the format eth0,1,eth1,2.
Task Customization	
Tasks	Displays the list of all tasks that will be executed as a part of the following application migration operations: ■ Setup Storage ■ Rehearse ■ Migrate
Add Pause	Allows you to add a pause after any predefined task.
Add Script	Allows you to add a task for custom script execution on a specific cluster node.
Delete	Allows you to delete a pause task or custom script execution task that you added. You cannot delete a predefined task.
Summary	Displays a summary of the application migration plan.

While creating the application migration plan, a query is executed on the source and target cluster nodes and this could cause a delay when moving from the Cluster Information & Dependency panel to the Target Disk Information panel. The query fetches the following details:

- VxVM Mirroring—Existing mirroring configuration
- VVR Replication:
  - Existing RVG configuration
  - Space validation on source disk group
  - Space required for target disk group
  - Network details like NIC, IP, and netmask of all the source and target cluster nodes

To create an application migration plan

1. In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane and click **Application Migration Plan**.
2. In the Migration Plan pane, right-click and click **Create**.

   The Create Migration Plan wizard opens.
3. In the Components Selection panel, specify a name and description, select the data migration type, and select a service group and target cluster.

   Click **Next**.
4. In the Cluster Information & Dependency panel, you can view the cluster details and modify the dependent service groups, if required.

   Use existing service group(s) in the target cluster and **Create new service group(s) in the target cluster** options are disabled and you will not be able to proceed if there are few service groups or resources in the target cluster which are present in the service group dependency tree that is being migrated. In this scenario, you have to cancel the operation, do one of the following, and then create the plan again:

   - Delete the conflicting service groups/resources on the target—This enables the **Create new service group(s) in the target cluster** option.

   Or

   - Create the cluster configuration in the target cluster manually—This enables the **Use existing service group(s) in the target cluster** option.

   Click **Next**.

   If you selected VxVM Mirroring as the Data Migration Type, go to 5.

   If you selected VVR Replication as the Data Migration Type, go to 6.

5 In the Target Disk Information panel, you can view the disk groups and select disks for creating mirror disk groups, if required.

If any disk group which is mirrored with the source disk group is configured under DiskGroup or CVMVolDg resource, this disk group is also displayed and is marked **Not Mirrored**. However, the source disk group is marked as **Mirrored**. In this scenario, you can ignore the mirror disk group and choose **Use Existing** or **Create New** for the source disk group.

If the disk group is mirrored and the disks used to create the mirror disk group are not present in all nodes of the target cluster, the disk group is marked as **Mirrored**. However, the **Use Existing** option will not be enabled.

If the disk group resources are not online in the source cluster or the XPRTLD daemon is not running in the source cluster nodes, the wizard will not be able to detect the state of the disk groups. In this scenario, the disk groups will not be marked as Mirrored or Not Mirrored.

If there are no disk groups for migration or if you select **Ignore** for all the disk groups listed, an error is displayed and you will not be able to proceed.

Click **Next**. Go to 7.

6 In the RVG Information Panel, you can view the disk groups, select disks, and enter RLINK details for creating replication configuration.

Disk groups are marked as **Replicated** if RVGs are configured in it; it does not depend on whether a secondary site is added or not to these RVGs. Disk groups are marked as **Not Replicated** if the disk groups does not have any RVGs configured.

If there are no disk groups for migration or if you select **Ignore** for all the disk groups listed, an error is displayed and you will not be able to proceed.

Click **Next**.

7 The Network Configuration panel displays details of network resources configured in any service groups being migrated. You can edit the target network details of these resources.

Note: In the Rehearse operation, IP-related resources are converted to FileOnOff type resources as part of cluster configuration in the target so that resource dependencies are maintained and to ensure that there will be no network connectivity to the applications which are brought online in the target, but NIC related resources are configured.

In Migrate operation, both IP and NIC related resources are configured in the target cluster.

For VVR Replication, IP resources configured to monitor IPs used for any existing replication configuration are not displayed.

Click **Next**.

8 In the Task Customization panel, you can view the list of all tasks that will be executed as a part of each operation. You cannot modify or delete a predefined task. After any predefined task in an operation, you can add a pause task or custom script execution task. These tasks are executed in the order specified during the course of the operation. You can also see whether a task is marked as critical or not. If any critical task fails, the entire operation fails.

See "Understanding user-defined tasks" on page 254.

Click **Next**.

9 The Summary panel provides a summary of the application migration plan. Review the plan and click **Finish** to create the application migration plan.

Click **Close** to close the wizard.

See "Executing the application migration plan" on page 270.

See "Understanding application migration operations" on page 255.

See "Understanding the cleanup operation" on page 259.

Understanding user-defined tasks

Pause task

During the application migration process, the operation can be paused for user intervention. When you create the plan, you can insert a pause task after any predefined task in an operation. When the plan is executed, the operation pauses

for an indefinite period at the point where the pause task was inserted. Click **Resume** to resume the operation.

Predefined pause tasks are added after the tasks listed below so that you can validate that task:

- Rehearse tasks—Cluster configuration, service group online, and service group offline.
- Migrate task—Cluster configuration.

Custom script execution task

You can execute custom scripts on cluster nodes during the application migration process. Scripts can be used to perform or validate application-specific tasks or any other tasks which you may want to perform during the operation. When you create the plan, you can add a custom script execution task after any predefined task in an operation. The script that will be executed should not accept any arguments. Any output from the script is logged in the operation execution log on the host where the script is executed.

You can mark a custom script execution task as critical or not. By default, the task is marked as non-critical. If the task is marked as critical, the return code obtained from the script after execution is used to determine the status of the task. If the return code is:

- Zero: Task succeeded
- Non-zero: Task failed

Understanding application migration operations

The following are the three main operations executed during application migration:

- Setup storage—See "Understanding the Setup Storage operation" on page 255.
- Rehearse—See "Understanding the Rehearse operation" on page 257.
- Migrate—See "Understanding the Migrate operation" on page 258.

Understanding the Setup Storage operation

In this operation, the storage tasks are taken care of.

Mirroring

If you chose to create mirror disk groups, the mirror disk groups are created and data sync is ensured.

The operation aborts if all volumes of all the disk groups being migrated is mirrored with the corresponding volumes in the mirror disk group. In this scenario, 100% sync between all volumes is not verified. This check is performed as part of the Rehearse or Migrate operation. If only some volumes are mirrored, the system is cleaned up so that the Setup Storage operation can be performed again. Only for the disk groups where the mirror was created by the application migration plan are cleaned up; mirrors created earlier are not cleaned up.

Replication

In this operation, Replication configuration is created and replication is established between the source cluster node and target cluster for disk groups marked as **Create New**. Initially, a basic verification is done to ensure that all disk groups marked **Use Existing** has RVGs configured and secondaries added and disk groups marked with **Create New** should have RVG configuration only if it was pre-created before creating the plan. All verifications are done in par with the configuration discovered during plan creation. Once verification is complete, DCM log is added to all volumes in the disk group after which SRL Volume is created as per the size calculations. Primary RVG is then created on the disk groups marked as **Create New**.

Migration plan creates a single RVG for each disk group and will have all volumes of the disk group made part of it. If RVGs are already pre-created on a disk group, these RVGs are used. Disk groups, volumes, and volume sets (if present) are then created on the target cluster node according to its size calculated from the source cluster node. After this is done, the target cluster node is added as a secondary site to the Primary RVG after which the replication is started.

Setup Storage operation completes only when Replication Sync is 100% complete. If any application writes are happening to the volumes on the source, it may take more time for the sync to complete. If you notice the sync is not progressing for a long time, check whether SRL has not overflowed in the logs and if so, it is recommended to reduce or stop application writes so that sync gets completed. Replication configuration created by the plan will be with default values and finetuning is also not performed.

On the source cluster, IP resources are added to the service group where the disk group resource is configured. These resources monitors the IPs used for replication. If the disk group is shared (CVM), IP resource and RVGLogowner resources are created in a separate failover service group with the name `amVvR_sg_failover_plan_name` and is linked to the service group which has the CVMVolDG resource configured.

On the target cluster, a service group is created by the name `amVvR_sg_plan_name`. IP and disk group resources for the entities created by the plan is added to this service group. In case of a shared disk group, a parallel service group by the name

amVvR_sg_parallel_*plan_name* is created in which CVMVolDG resource is added and another failover service group is created amVvR_sg_failover_*plan_name* which contains the IP and RVGLogowner resources. The operation aborts if all RVGs in all disk groups which are part of the migration have secondaries added to it.

Understanding the Rehearse operation

In this operation, you can bring application online on the target cluster and test the application before performing the actual migration.

Mirroring

In this operation, the sync status of all volumes are checked after which cluster configuration of the selected service group is discovered in the source and translated to the target. The mirror disk group is then detached from the source cluster nodes and endian changes are performed on all volumes of the mirror disk group in the target cluster nodes.

After the endian changes are done, the service groups are brought online in the target cluster. After ensuring the application is running fine in the target, the service groups are taken offline and the target cluster configuration is removed.

Before removing the cluster configuration, a backup of the configuration is taken on the first node of the target cluster in the /etc/VRTSvcs/conf/config directory and the name will be in the following format:

main.cf_*plan_name.date.time*

The volumes in the mirror disk groups are then reattached to the corresponding volumes in the source disk group.

The operation aborts if all volumes between the source and mirror disk groups are not completely synced.

Replication

Initially, all RVGs are checked to ensure 100% sync. At times, if the application is writing to the volumes, data sync might be in progress to the secondary site and hence the Rehearse operation will not proceed. In such scenario, you can reduce or stop application writes so that data remains in sync.

Pre-requisites are done on target cluster nodes in order to create space optimized snapshot of volumes such as preparing volumes, cache volume, and cache object creation. IP, DiskGroup/CVmVolDg, RVGLogowner resources are then removed from the target cluster for the disk groups which are being migrated as part of the plan. This is to ensure no duplicate resources appear for an entity during cluster translation.

Cluster configuration of the selected service group and its dependencies are then discovered on the source and translated to the target. The file systems on all mounted volumes of the disk groups which are part of the replication is then frozen on the source for a moment and space optimized snapshots of these volumes are taken on the target with the help of VVR In-Band Control Messaging (vxibc). After the snapshots are taken, endian changes are performed on these snapshots and the service groups are brought online on the target. When service groups are brought online, the snapshot volumes gets mounted on the target cluster. Applications started on the target can write on these snapshot volumes until the cache volume fills up. After ensuring the application is running fine on the target, the service groups are taken offline and the target cluster configuration is removed.

Before removing the cluster configuration, the configuration is backed up on the first node of the target cluster in the `/etc/VRTSvcs/conf/config` directory and the name is in the following format:

`main.cf_plan_name.date.time`

The disk groups are then imported on the target cluster node and the snapshots are destroyed. The IP, DiskGroup/CVMVolDg, and RVGLogowner resources are re-created on the target, as required, and then brought online.

Understanding the Migrate operation

In this operation, the actual migration happens. Service groups on the source are taken offline, so there will be a downtime when the operation is in progress. After the operation is complete, you can use the target cluster to run the application.

Mirroring

In this operation, the service groups in the source are taken offline before detaching the mirror disk groups from the source cluster nodes. This is to ensure data integrity before the final migration. After the mirror disk groups are detached from the source cluster nodes, endian changes are performed on these volumes in the target cluster nodes and the service groups are brought online in the target cluster.

The operation aborts if all volumes between the source and mirror disk groups are not completely synced.

Replication

In this operation, verification is done to ensure 100% sync between the source and target clusters. Cluster configuration is completely translated and service groups are taken offline from the source. Replication is stopped for all RVGs of all disk groups part of the migration. Endian changes are performed on volumes part of replication on the target after which service groups are brought online. Secondary

site is removed from the primary RVG and the primary RVG is also removed for all the disk groups marked as **Create New**.

Understanding the cleanup operation

Cleanup operation is an internal operation which takes place whenever a critical task fails in an operation. This operation takes care of cleaning up the managed hosts to ensure the system is in a clean and known state so that the operation can be performed again after taking corrective measures.

The status or the progress of this operation is not displayed in the user interface. To view the status of this operation, check the logs available in the managed hosts.

Mirroring: Steps performed as a part of cleanup:

- Setup Storage—If you had selected the **Create New** option in the wizard, mirror is removed for all volumes between source and target disk groups. Mirror disk groups are deleted.
- Rehearse and Migrate—All volumes of the mirror disk groups are reattached with the corresponding volumes in the source disk group. If the plan created the target cluster configuration, the configuration is backed up and then removed.

Replication: Steps performed as a part of cleanup:

- Setup Storage—Cleanup is performed only for disk groups marked as **Create New**. For these disk groups, replication is stopped (if established), secondary site is deleted and primary RVG is removed if the RVG was created by the plan. Disk groups created on the target cluster by the plan is also removed.
- Rehearse—Snapshots created for volumes are removed from the secondary. Cache object and cache volume are deleted. VCS resources for disk groups, IPs, etc. are restored. If the plan created the target cluster configuration, the configuration is backed up and then removed.
- Migrate—Replication, if stopped, is started again. If the plan created the target cluster configuration, the configuration is backed up and then removed.

Understanding the tasks executed in each operation

Endian changes

The application peration uses the `fscdsconv` utility to perform endian changes on the target cluster node. This utility creates a recovery file for the purpose of

conversion. The default path for the creation of the recovery file is: `/var/opt/VRTSsfmh/tmp`. The file system `/var` on the target cluster node must have more than 5 GB space.

If the file system does not have sufficient space for creating the file

1. On all nodes of the target cluster, create the `/opt/VRTSsfmh/etc/appmig.conf` file.
2. In the `appmig.conf` file, define the APPMIG_RECOVERY_PATH=*path* variable. Here, *path* is the full directory path where the recovery file needs to be created.

To check how much space is required in a VxFS file system for the recovery file as used by the `fscdsconv` utility, see: http://www.veritas.com/docs/000012748

Cluster configuration translation at the service group level

The following points describe how cluster configuration is translated at the service group level:

- All the nodes in the target cluster would be part of the SystemList for the service group in the target.
- Service groups on the source cluster that is being migrated are translated to the target cluster along with dependencies.
- While migrating service groups between CVM clusters, the CVM service group is not translated, but its dependencies are translated. To ensue that the dependencies are translated, Veritas recommends that you should not deselect the CVM service group during plan creation.
- Any setting with the following service group level attributes will be ignored in the target cluster:
  - SystemZones
  - VCSi3Info
  - TriggerPath
  - ClusterList
  - Frozen
  - TriggersEnabled
  - TypeDependencies
  - AdministratorGroups
  - OperatorGroups
  - Administrators

  - Operators
  - Guests
- Triggers configured in the source cluster will not be translated to the target cluster. If you require triggers, trigger files must be manually copied and trigger attributes must be manually configured in the target cluster. You can use the pause task after the **Configure Cluster on Target** task or you can create a custom script to automate this.
- If the source and target operating systems are different, the ContainerInfo attribute at the group level is skipped in the target cluster.
- If a service group attribute is localized in the source cluster and the number of nodes in the target cluster is more than the number of nodes in the source cluster SystemList, the localized attribute is converted to a global attribute.
- If the source cluster has offline group dependencies and the target cluster is a single node cluster, offline group dependencies are deleted from the target cluster. Remote dependencies are converted to local dependencies.
- All site dependencies are deleted from the target cluster.

Cluster configuration translation at the resource level

The following points describe how cluster configuration is translated at the resource level:

- Resource attributes part of the service groups being migrated are translated from the source to target cluster. For example, if the DiskGroup attribute exists in the service group, the attribute value is translated from the source cluster to the target cluster.
- VxVM Mirroring—For each disk group resource, another disk group resource is created by the application migration operation. The name of this resource will start with `amMir_`. This is for internal housekeeping of snap disk group created for mirroring. This resource is skipped in the target cluster.
- If the source and target operating systems are different, ResContainerInfo and ContainerOpts attributes are skipped in the target cluster.
- If a resource attribute is localized in the source cluster and the number of nodes in the target cluster is more than the number of nodes in the source cluster SystemList, the localized attribute is converted to a global attribute.
- During Rehearse operation, IP resources like IP, IPMultiNIC, and IPMultiNICB are not created in the target cluster. This is to avoid accidental connection of clients to the application during Rehearse operation. These resources are converted to FileOnOff type to maintain resource dependencies.

- If any resource type configured in the source cluster is not available in target cluster, the corresponding resources are removed from the target cluster during cluster configuration.
- If the source and target cluster operating systems are different, IPMultiNICB resources are converted to IP type and MultiNICB type resources are converted to NIC type.
- VVR Replication—In the Rehearse operation, Volume and VolumeSet type resources managing volumes and volume sets being replicated are converted to FileOnOff type resource in the target cluster. This is done as snapshot volumes created in the target cluster cannot be managed using these resources.
- VVR Replication—In the Migrate operation, Volume and VolumeSet type resources managing volumes and volume sets being migrated are converted to FileOnOff type in the source cluster. This is to ensure taking the service group offline does not fail in the source cluster.

List of tasks executed

Table 21-4 lists predefined tasks executed as a part of each operation. Each predefined task in an operation is marked as **Critical** or not. The table also lists the tasks for which the cleanup operation is performed and the state of the managed hosts after the operation.

If a task marked as critical fails, the operation aborts. If a task is not critical, the operation continues even if it fails. For most critical tasks, failure necessitates a cleanup of the system and an internal cleanup operation is performed.

See “Understanding the cleanup operation” on page 259.

Table 21-4 Mirroring—List of tasks executed

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Setup Storage			
Deploy scripts on the managed host	Yes	No	No change.
Verify mirror configuration	Yes	No	No change.
Prepare disk group for mirroring	Yes	Yes	After cleanup, system reverts to the state before the Setup Storage operation. Perform the Setup Storage operation again.

Table 21-4 Mirroring—List of tasks executed *(continued)*

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Create mirror between source disk group	Yes	Yes	After cleanup, system reverts to the state before the Setup Storage operation. Perform the Setup Storage operation again.
Snapshot sync progress	Yes	No	System will be in existing state and mirror will be under sync on the managed host.
Remove scripts from the managed host	No	No	System will be in existing state.
Rehearse			
Deploy scripts on the managed host	Yes	No	No change.
Verify mirror configuration	Yes	No	No change.
Configure target cluster	No	No	Target cluster configuration is created partially. You must validate the configuration and make necessary changes during the Pause task before proceeding.
Pause plan for user validation	Yes	No	No change.
Detach mirror disk group	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, mirror is reattached to the source. Target cluster configuration is removed after taking backup. Perform the Rehearse operation again.
Online service group in the target cluster	No	No	Service groups will be in partial state. Check the reason in the Pause task.
Pause plan for user validation	Yes	No	No change.
Offline service group in the target cluster	No	No	Service groups will be in partial state. Take all the service groups offline and check whether the mirror disk groups are deported; if not, manually deport the mirror disk groups in the Pause task.
Pause plan for user validation	Yes	No	No change.

Table 21-4 Mirroring—List of tasks executed *(continued)*

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Remove service group from the target cluster	Yes	No	Task will not fail.
Endian changes on the mirror disk group	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, mirror is reattached to the source. Perform the Rehearse operation again.
Reattach mirror disk group in the source cluster	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, mirror is reattached to the source. Perform the Rehearse operation again.
Snapshot sync progress	Yes	No	System will be in existing state. Mirror will be under sync on the managed host.
Remove scripts from the managed host	No	No	System will be in existing state.
Migrate			
Deploy scripts on the managed host	Yes	No	No change.
Verify snap mirror configuration	Yes	No	No change.
Configure target cluster	No	No	Target cluster configuration is partially created. You must validate the configuration and make necessary changes before proceeding.
Pause plan for user validation	Yes	No	No change.
Offline service group in the source cluster	Yes	No	Service groups will be in partial state. Do the following: ■ Check the reason why taking the service group offline failed. ■ Fix the issue. ■ Bring the service groups online. ■ Start the Migrate operation.
Detach snap mirror disk group	Yes	Yes	During cleanup, mirror is reattached to the source and target cluster configuration is removed. Bring the service groups online in the source cluster and perform the Migrate operation again.

Table 21-4 Mirroring—List of tasks executed *(continued)*

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Endian changes on the mirror disk group	Yes	Yes	During cleanup, mirror is reattached to the source and target cluster configuration is removed. Bring the service groups online in the source cluster and perform the Migrate operation again.
Online service group in the target cluster	No	No	Service groups will be in partial state in the target cluster.
Remove scripts from the managed host	No	No	System will be in existing state.

Table 21-5 Replication—Lists of tasks executed

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Setup Storage			
Deploy scripts on managed hosts	Yes	No	No change.
Verify replication configuration	Yes	No	No change.
Configure replication on source and target clusters	Yes	Yes	After cleanup, system reverts to the state before the Setup Storage operation. Check and resolve the issue. Perform the Setup Storage operation again.
Set up replication between source and target clusters	Yes	Yes	After cleanup, system reverts to the state before the Setup Storage operation. Check and resolve the issue. Perform the Setup Storage operation again.
Create RVGLogowner type resource for shared disk group(s)	Yes	Yes	Task will not fail.
Replication sync progress	Yes	No	System will be in existing state and replication will be under sync between source and target cluster nodes.
Remove scripts from managed hosts	No	No	System will be in existing state.
Rehearse			

Table 21-5 Replication—Lists of tasks executed *(continued)*

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Deploy scripts on managed hosts	Yes	No	No change.
Verify replication configuration	Yes	No	No change.
Delete replication-related VCS resources on the target cluster	Yes	Yes	Task will not fail.
Perform target cluster configuration	No	No	Target cluster configuration is created partially. You must validate the configuration and make necessary changes during the Pause task before proceeding.
Pause plan for user validation	Yes	No	No change.
Set up volumes on target cluster for creating snapshot	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, snapshots are deleted. Target cluster configuration is removed. Replication-related resources are re-created on the target cluster. Check and resolve the issue. Restart the Rehearse operation.
Create snapshots on target by freezing VxFS mount points and replication	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, snapshots are deleted. Target cluster configuration is removed. Replication-related resources are re-created on the target cluster. Check and resolve the issue. Restart the Rehearse operation.
Perform endian changes on snapshot volumes on the target cluster	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, snapshots are deleted. Target cluster configuration is removed. Replication-related resources are re-created on the target cluster. Check and resolve the issue. Restart the Rehearse operation.
Bring service groups online on the target cluster	No	No	Service groups will be in partial state. Check the reason for the failure and correct it.
Pause plan for user validation	Yes	No	No change.

Table 21-5 Replication—Lists of tasks executed *(continued)*

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Take service groups offline on the target cluster	No	No	Service groups will be in partial state. Take all the service groups offline and during the Pause task, check whether all the disk groups are deported.
Pause plan for user validation	Yes	No	No change.
Remove target cluster configuration	Yes	Yes	Task will not fail.
Remove snapshot volumes and re-create replication-related VCS resources on the target cluster	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, snapshot removal is reattempted and replication-related resources are re-created on the target cluster.
Remove scripts from managed hosts	No	No	No change.
Migrate			
Deploy scripts on managed hosts	Yes	No	No change.
Verify replication configuration	Yes	No	No change.
Delete replication-related VCS resources on the target cluster	Yes	Yes	Task will not fail.
Perform target cluster configuration	No	No	Target cluster configuration is created partially. You must validate the configuration and make necessary changes during the Pause task before proceeding.
Pause plan for user validation	Yes	No	No change.
Reconfigure the source cluster to prevent plan failure	Yes	Yes	During cleanup, cluster configuration performed on the target is removed. Restart the Migrate operation

Table 21-5 Replication—Lists of tasks executed *(continued)*

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Take service groups offline on the source cluster	Yes	No	Service groups will be in partial state. Do the following: ■ Check the reason why taking the service group offline failed. ■ Fix the issue. ■ Bring the service groups online. ■ Start the Migrate operation.
Stop replication and remove the secondary site	Yes	Yes	During cleanup, secondary sites which might have been removed will be re-added and replication is started. Service groups will be in offline state on the source cluster. Bring the service groups back online. Check and resolve the issue. Restart the Migrate operation.
Perform endian changes on replicated volumes	Yes	Yes	During cleanup, secondary sites which might have been removed will be re-added and replication is started. Service groups will be in offline state on the source cluster. Bring the service groups back online. Check and resolve the issue. Restart the Migrate operation.
Bring service group online on the target cluster	No	No	Service groups will be in partial state on the target cluster.
Remove replication configuration from source and target cluster	No	No	Replication configurations added on the source and target cluster disk groups will not be fully cleaned up.
Remove scripts from managed hosts	No	No	No change.

Validations performed before migration plan execution

Before a plan is executed, validations are performed in order to check the state of the clusters, service groups, and the resources monitoring the disk groups.

In some scenarios:

- If the validation fails, the plan will be marked as an **Invalid Plan**.
- Even though the plan is marked as Invalid, the plan can be executed after taking necessary corrective steps.

- After taking the corrective steps, you must edit the plan before executing it.

Table 21-6 lists the validations performed before a plan is executed and the state of the plan if the validation fails. It also lists when the plan will be marked as invalid or not and the necessary corrective steps.

Table 21-6 Validations performed

Object	State	Is the plan invalid?	Will the plan remain invalid forever?	Corrective Action
Source cluster	Deleted	Yes	Yes	Edit the plan, select any service group of an existing source cluster that needs to be migrated and execute the plan again.
Source cluster	Faulted	No	No	Resolve the issues in the source cluster. Ensure that the cluster is not in faulted state in VIOM and execute the plan.
Target cluster	Deleted	Yes	Yes	Edit the plan, select an existing target cluster to which migration needs to be performed and execute the plan.
Target cluster	Faulted	No	No	Resolve the issues in the target cluster. Ensure that the cluster is not in faulted state in VIOM and execute the plan.
Service group	Deleted	Yes	Yes	Edit the plan and reselect any existing service group in the source cluster that needs to be migrated and execute the plan.
Dependent service groups	Deleted	Yes	No	Add the dependent service group that was deleted and execute the plan.
Disk group resource	Deleted	Yes	No	Add the disk group resource that was deleted and execute the plan.
Disk group resource offline	Offline	No	No	Bring the disk group resource online and execute the plan.

Table 21-6 Validations performed *(continued)*

Object	State	Is the plan invalid?	Will the plan remain invalid forever?	Corrective Action
Disk group	Deleted	Yes	Yes	Edit the plan so that the current state of the managed hosts is detected by the wizard. Save the plan and execute it.
Volume	Deleted	Yes	Yes	Edit the plan so that the current state of the managed hosts is detected by the wizard. Save the plan and execute it.

Executing the application migration plan

After creating the migration plan, you must execute the plan.

If you had selected VVR Replication as the data migration type, you must have a valid VVR license and you must start the following services on the source and target cluster nodes before executing the plan:

- vxnm-vxnetd
- vras-vradmind
- vxrsyncd

To execute the application migration plan

1. In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane and click **Application Migration Plan**.
2. Right-click the plan that you created, click **Setup Storage**, and click **OK** in the Execute Migration Plan window.

   On the **Latest Run** area of the screen, you can view the Setup Storage tasks that are being executed and the corresponding status.

   The **Status** column of the Migration Plan window displays the overall status of the Setup Storage operation.

   To verify whether mirror disk groups have been created in the source cluster in case of mirroring and whether remote disk groups have been created in the target cluster in case of replication:

   - In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

- Expand the Organization, or **Uncategorized Hosts** to locate the managed host.
- Expand the managed host of the cluster and then expand **Disk Groups**.

To verify whether RVGs have been created in the disk group in case of replication:

- Expand the managed host in the target cluster.
- Expand RVGs.

If you had chose to pause any task as outlined in step 8 of Creating an application migration plan section, click **Resume** to resume the task that was paused.

3 Go to step 1. Right-click the plan that you created, click **Rehearse**, and click **OK** in the Execute Migration Plan window. You can execute the Rehearse operation multiple times.

On the **Latest Run** area of the screen, you can view the Rehearse tasks that are being executed and the corresponding status.

The **Status** column of the Migration Plan window displays the overall status of the Rehearse operation.

If you had chose to pause any task as outlined in step 8 of Creating an application migration plan section, click **Resume** to resume the task that was paused.

4 Go to step 1. Right-click the plan that you created, click **Migrate**, and click **OK** in the Execute Migration Plan window.

You can execute the Migrate operation only once. After you start the Migrate operation, you cannot stop or revert it. You must perform this operation only after executing all the other operations.

On the **Latest Run** area of the screen, you can view the Migrate tasks that are being executed and the corresponding status.

The **Status** column of the Migration Plan window displays the overall status of the Migrate operation.

If you had chose to pause any task as outlined in step 8 of Creating an application migration plan section, click **Resume** to resume the task that was paused.

5 To verify whether the service groups have been migrated successfully to the target cluster:

- In the Management Server console, go to the **Availability** perspective and expand **Manage** in the left pane.

- Expand the Organization, or **Uncategorized Clusters** to locate the cluster.
- Go to the source cluster and expand **Service Groups**, the service groups must be offline.
- Go to the target cluster and expand **Service Groups**, the service groups must be online.

Refer to the following sections:

See "Understanding the tasks executed in each operation" on page 259.

See "Validations performed before migration plan execution" on page 268.

Editing an application migration plan

After you create a plan, you can edit it if you need to modify any of the entries. When you edit a plan, the service group dependencies and disk group details are not populated from the previously saved plan, but are fetched afresh. This is to ensure that the current state of the service groups and disk groups in the source cluster are populated when you edit the plan.

Before you edit a plan, note the following:

- If you edit the plan before the Setup Storage operation:
  Mirroring—Re-select the disks for any mirror disk group which needs to be created for each source disk group.
  Replication—Re-select the disks and re-enter the RLINK details.
- Mirroring—If you edit a plan after the Setup Storage or Rehearse operation, mirror disk groups are also be displayed in the **Target Disk Information** panel. You must select the **Use Existing** option for source disk groups and select **Ignore** for mirror disk groups, if these were created as part of Setup Storage. The mirror disk groups will have a naming convention as `amMir_source_diskgroup_name`.
- If you edit the plan after it is imported, any custom script execution task which was part of the imported plan must be removed and new custom script execution tasks must be added as necessary by uploading the custom scripts copied when the plan was exported.

To edit an application migration plan

1. In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane, and click **Application Migration Plan**.
2. In the Migration Plan pane, right-click and click **Edit**.

   The Edit Migration Plan wizard opens.

3 Edit the fields in the migration plan as necessary.

4 Click **Save** to the save the changes in the existing plan or click **Save As** to save the changes as a new plan.

Note: If you want to use the Save As option, you must change the plan name.

Deleting application migration plan(s)

You can delete a single plan or multiple plans. To select multiple plans, hold the **Ctrl** key and select the plans.

To delete application migration plan(s)

1 In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane, and click **Application Migration Plan**.

2 In the Migration Plan pane, select the plan(s), right-click, and select **Delete**.

3 In the Delete Migration Plan(s) window, click **OK**.

The selected plan(s) are deleted.

Exporting application migration plan(s)

You can export any application migration plan to your system. It can be a single plan or multiple plans. Exported plans are saved in JavaScript Object Notation (JSON) format. Use a JSON viewer to view the contents of the exported file. The file created upon export can be used for importing the plan, if a plan needs to be imported on same or a different Management Server.

When you export multiple plans simultaneously, only a single JSON file is created. This file contains all the entries of the exported plans. You can edit the JSON file, but you must not modify the JSON syntax structure of the file. If you modify the JSON syntax structure, the import operation fails. If multiple plans are part of a single JSON file, you can remove the entries of a plan from the JSON file without modifying the syntax structure. Veritas recommends to change the value of only the following entries in the JSON file. Based on the value of these entries, the value of the remaining entities are fetched when the plan is being imported:

- planName—Name of the plan
- description—Description of the plan
- selectedSGName—Name of the selected service group

- sourceClusName—Name of the source cluster
- targetClusName—Name of the target cluster

When you edit the plan after exporting, you need to edit only the names of the entities listed above; the IDs are fetched during the import operation. If the entities mentioned are not available on the Management Server where the plan is being imported, the import operation fails.

Note: Custom scripts part of the plan are not exported. You have to manually copy the custom scripts either from the Management Server (path is available in the exported JSON file) or from the location from which you uploaded the custom scripts and use it when you edit the plan.

To export application migration plan(s)

1. In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane, and click **Application Migration Plan**.
2. In the Migration Plan pane, select a plan, right-click, and click **Export**.

   The Export Plan(s) window opens.
3. Click **Yes**.

   The selected plan(s) are exported.

See "Importing application migration plan(s)" on page 274.

Importing application migration plan(s)

You can import any application migration plan from your system. It can be a single plan or multiple plans. To import multiple plans, all the plans must be part of the same file.

The following validations are performed during the import operation. If any of these validations fail, the plan will not be imported.

- No plan with same name as the plan being imported must exist.
- Plan name must be according to the recommendation in the Creating an application migration plan section.
- The names of the source cluster, target cluster, and the selected service group mentioned in the JSON file used for import must be available on the Management Server in which the plan is being imported.

If the plan is successfully imported, the plan description is prefixed with the [Imported] tag.

After the plan is imported, Veritas recommends to edit the plan to validate all the entries. If any custom script execution tasks were part of the plan, these entries are displayed in the Task Customization panel of the wizard. As the custom scripts may not be available in the Management Server where the plan is imported, delete these pre-existing tasks and add new custom script execution tasks in the same place by uploading custom scripts copied during the export operation.

To import application migration plan(s)

1. In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane, and click **Application Migration Plan**.
2. In the Migration Plan pane, right-click and click **Import**.

   The Import Plan(s) wizard opens.
3. Click **Browse** to select the plan(s).
4. Click **Import**.

   The selected plans are imported.

See "Exporting application migration plan(s)" on page 273.

Viewing historical runs

The Historical Runs option allows you to view the execution history of a plan.

To view historical runs

1. In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane, and click **Application Migration Plan**.
2. In the Migration Plan pane, select a plan, right-click and click **Historical Runs**.

   The Historical Runs window displays the historical runs of the selected plan.

Viewing properties of an application migration plan

To view the properties of an application migration plan

1. In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane, and click **Application Migration Plan**.
2. In the Migration Plan pane, select a plan, right-click and click **Properties**.

   The Properties window displays the following:

   - Summary of the plan

- Disk group details
- Custom scripts and its location on the Management Server
- Network details
- RVG details (Displayed only if the Data Migration Type is VVR Replication)

Application migration logs

The following table lists the name and location of the log files. You can use these logs for troubleshooting.

Table 21-7 Log files

Name	Path
Add-on installation and uninstallation related logs	■ Windows MS: C:\ProgramData\Symantec\VRTSsfmcs\logs\vcsappmig_deployment.log ■ Linux MS: /var/opt/VRTSsfmcs/logs/vcsappmig_deployment.log
Application migration wizard logs	■ Windows MS: C:\ProgramData\Symantec\VRTSsfmcs\logs\WebDebugLog.txt C:\ProgramData\Symantec\VRTSsfmcs\logs\tomcat.log ■ Linux MS: /var/opt/VRTSsfmcs/logs/WebDebugLog.txt /var/opt/VRTSsfmcs/logs/tomcat.log
Application migration execution logs	AIX, Linux, and Solaris managed hosts: /var/opt/VRTSsfmh/logs/app_migration_*plan_name*.log

Index

A

additional scripts
- about 180
- check_env.pl 180
- readme.html 180

array port
- performance graphs 103

B

Brocade switch discovery 34
- about 28
- prerequisites 28

C

Cisco switch discovery 31
- about 28
- prerequisites 28

classifications for LUNs
- creating definitions 107
- deleting definitions 109
- modifying definitions 108
- modifying the order 110
- overview 106
- refreshing information 110

controller performance graphs 104
creating storage templates using NTFS file systems 205
creating storage templates using volumes 206
creating storage templates using VxFS file systems 199
custom scripts for SFHA hot fix deployment
- adding to hot fixes 46
- removing or replacing 47
- requirements 44

customized solutions 21

D

deep array discovery
- 3PAR 79
- editing configuration 81

deep array discovery *(continued)*
- EMC Celerra enclosures 75
- EMC CLARiiON enclosures 72
- EMC Symmetrix enclosures 68
- EMC VNX enclosure 76
- EMC VPLEX enclosures 78
- HITACHI enclosures 54
- HP EVA enclosures 73
- IBM SVC enclosures 80
- IBM System Storage DS 74
- IBM XIV enclosures 69
- NetApp enclosures 70
- refresh enclosures 95
- removing configuration 94

deleting storage templates 210
deploying customized solutions 20
device identifier 181
- enclosure identifier formula 190
- logical disk identifier formula 190
- SCSI Page 190
- SCSI Product 190
- SCSI Vendor 190

device identifier formulas
- Field Format 181
- No of Bytes 181
- Page 83 Id Type 181
- SCSI Page No 181
- Start Byte 181

disabling performance metering for enclosure 101
discovery command output
- --list adapters --encl 167
- --list capacities --encl 168
- --list encls 162
- --list ldevhostmap --encl 171
- --list ldevpdevmap --encl 170
- --list ldevs --encl 165
- --list meta-ldevs --encl 172
- --list pdevs --encl 164
- --list ports --encl 167
- --list raidgroups --encl 172
- --list replications --encl 177
- --list rgldevmap --encl 175

discovery command output *(continued)*
--list rgpdevmap --encl 175
--list thinpools --encl 173
--list tpldevmap --encl 177
--list tpsrcldevmap --encl 176
discovery host 53
discovery script
about 157
arguments 158
parameters 157
Distribution Manager Add-on
about 20
downloading storage templates 209

E

enabling
enclosure performance metering 101
enclosure capabilities
logical characteristics 186
physical characteristics 186
enclosure performance metering
disabling 101
enabling 101

F

fabric discovery
about 26

H

hot fix deployment 43
custom scripts 44, 46–47
prerequisites 42

I

Impact Analysis report for volume migration 213

L

locking storage templates 210
LUNs
creating classification definitions 107
deleting classification definitions 109
modifying classification definitions 108
modifying classification levels 110
overview of classifications 106

M

migrating volumes 212
moving disk groups
between enclosures 219
moving volumes
between enclosures 214, 217

N

network-attached storage (NAS) discovery 54

O

overview
Distribution Manager Add-on 20
Patch Installer Add-on 40
Storage Provisioning and Enclosure Migration Add-on 197
Veritas Operations Manager Storage Insight Add-on for Deep Array Discovery and Mapping 51
Veritas Operations Manager Storage Insight SDK Add-on 152

P

Patch Installer Add-on
about 40
adding custom scripts to hot fixes 46
deploying hot fixes 43
prerequisites 42
performance graphs
array port 103
controller 104
enclosure 102
performance metering statistics 99
plug-in vitals
additional scripts 184
discovery method 184
discovery script 184
provisioning storage 208

R

refreshing classification information 110

S

scripts for SFHA hot fix deployment. *See* custom scripts for SFHA hot fix deployment
seed switch 27
SFHA hot fix deployment 43
custom scripts 44, 46–47
prerequisites 42

SNMPv1/2c
Community String 32
SNMPv3
AuthNoPriv 32
AuthPriv 32
NoAuthNoPriv 32
storage enclosures
discovered information 52
discovery methods 52
Storage Insight plug-in
create 183
device identifier formulas 181
discovery script arguments 158
edit 192
test with collected data 193
test with live data 193
Storage Insight plug-in sample
check_env.pl 182
compellent_data.tar.gz 182
dellcompcli.ps1 182
Readme.html 182
SI_DELL_COMPELLENT_POWERSHELL-1.0.0.sfa 182
Storage Insight SDK Add-on
about 152
discovered objects 153
discovery command output 161
installing 157
objects discovery order 160
SI_PLUGIN_DIR 160
SI_PLUGIN_WORK_DIR 160
storage provisioning 208
overview 197
Storage Provisioning and Enclosure Migration Add-on
about 197
storage template
creating using NTFS file systems 205
creating using volumes 206
creating using VxFS file systems 199
deleting 210
downloading 209
locking 210
unlocking 210
updating 207
uploading 209
using to provision storage 208
switch discovery
edit 38
refresh 37
remove 38

T

template. *See* storage template
tiers for LUNs. *See* classifications for LUNs
troubleshooting
Brocade switch discovery 29
Cisco switch discovery 29

U

unlocking storage templates 210
updating storage templates 207
uploading storage templates 209

V

VCS clusters
password change 17
Veritas Cluster Server Utilities Add-on
features 15
limitations 15
Overview 14
prerequisites 16
Veritas HA Plug-in for vSphere Web Client
about 223
limitations 224
Management Server High Availability environment 230
register HA Plug-in 228
unregister HA Plug-in 229
volume migration
about 212
by disk group 219
by enclosure 217
by host 214
deleting from schedule 221
impact analysis 213
pausing or resuming 221