

IBM PowerHA SystemMirror for AIX

Standard Edition

Version 7.2

*PowerHA SystemMirror
for IBM Systems Director*

IBM

IBM PowerHA SystemMirror for AIX

Standard Edition

Version 7.2

*PowerHA SystemMirror
for IBM Systems Director*

IBM

Note

Before using this information and the product it supports, read the information in "Notices" on page 179.

This edition applies to IBM PowerHA SystemMirror 7.2 Standard Edition for AIX and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2015.**

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

Contents

About this document	v	lsdependency command	55
Highlighting	v	lsefsKeystore command	57
ISO 9000.	v	lsfallbacktimer command	58
Related information	v	lsfilecollection command	59
PowerHA SystemMirror for IBM Systems Director	1	lsgroup command	60
What's new in for PowerHA SystemMirror for IBM Systems Director	1	lsinterface command	61
Installing PowerHA SystemMirror for IBM Systems Director	4	lsldapcl command	63
Prerequisites for installing PowerHA SystemMirror for IBM Systems Director	5	lsldapsrv command.	64
Installing PowerHA SystemMirror for IBM Systems Director on the server	5	lslog command	65
Installing PowerHA SystemMirror for IBM Systems Director on nodes.	6	lsmethod command	66
Unlocking nodes in IBM Systems Director after installation has completed	7	lsnode command	67
Logging in to the IBM Systems Director Console ..	7	lsnetwork command	68
Navigating the PowerHA SystemMirror for IBM Systems Director interface	7	lspersistentip command	70
PowerHA SystemMirror for IBM Systems Director cluster simulator mode	9	lsresgrp command	71
Using offline simulation mode	9	lsserviceip command	72
Using planning mode	10	lssite command	73
Configuring clusters and resource groups by using PowerHA SystemMirror for IBM Systems Director	11	lssnapshot command	74
Configuring clusters using PowerHA SystemMirror for IBM Systems Director	11	lstape command	75
Configuring resource groups using PowerHA SystemMirror for IBM Systems Director	25	lsversion command.	77
Restoring a cluster by using snapshot	35	lsvolumegroup command.	78
Managing clusters and resource groups by using PowerHA SystemMirror for IBM Systems Director	36	mkappctl command	79
Managing clusters by using PowerHA SystemMirror for IBM Systems Director	36	mkappmon command	81
Managing resource groups by using PowerHA SystemMirror for IBM Systems Director	40	mkcluster command	84
Managing sites by using PowerHA SystemMirror for IBM Systems Director.	41	mkdependency command	87
Managing events by using PowerHA SystemMirror for IBM Systems Director	42	mkefsKeystore command	90
Managing custom verification methods by using PowerHA SystemMirror for IBM Systems Director	45	mkfallbacktimer command	91
Managing replicated storage.	46	mkfilecollection command	93
PowerHA SystemMirror for IBM Systems Director commands.	49	mkgroup command.	95
help command	49	mkinterface command	97
importcluster command	49	mkldapcl command	98
lsappctl command	51	mkldapsrv command.	100
lsappmon command	52	moveresgrp command	102
lsclrole command	53	mknetwork command	104
lscluster command	54	mknode command	105
		mkmethod command.	107
		mkresgrp command	109
		mksecurity command.	113
		mkserviceip command	115
		mksite command	116
		mksnapshot command	118
		mktape command	119
		mkuser command	121
		movresgrp command.	126
		rfcluster command	128
		rmappctl command	129
		rmappmon command	130
		rmcluster command	132
		rmdependency command	133
		rmefsKeystore command	135
		rmfallbacktimer command	136
		rmfilecollection command	137
		rmgroup command	139
		rminterface command	140
		rmldapcl command	141
		rmldapsrv command	142
		rmnetwork command	143

rmmethod command	145
rmnode command	146
rmpersistentip command	148
rmresgrp command	149
rmserviceip command	151
rmsite command	152
rmsnapshot command	154
rmtape command	155
rmuser command	156
rssnapshot command	158
startcluster command	159
startnode command	161
startresgrp command	163
stopcluster command	164
stopnode command	166

stopresgrp command	167
synccluster command	168
syncfilecollection command	170
undochanges command	171
vercluster command	172
vlog command	174
vreport command	176

Notices 179

Privacy policy considerations 181

Trademarks 181

Index 183

About this document

PowerHA® SystemMirror® for IBM® Systems Director is a plug-in for IBM Systems Director that uses a web browser interface to manage PowerHA SystemMirror capabilities.

Highlighting

The following highlighting conventions are used in this document:

Bold	Identifies commands, subroutines, keywords, files, structures, directories, and other items whose names are predefined by the system. Also identifies graphical objects such as buttons, labels, and icons that the user selects.
<i>Italics</i>	Identifies parameters whose actual names or values are to be supplied by the user.
Monospace	Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of portions of program code similar to what you might write as a programmer, messages from the system, or information you should actually type.

ISO 9000

ISO 9000 registered quality systems were used in the development and manufacturing of this product.

Related information

- The PowerHA SystemMirror PDF documents are available in the IBM Knowledge Center.
- The following PowerHA SystemMirror release notes are located in the following locations:
 - PowerHA SystemMirror Standard Edition for AIX®: `/usr/es/sbin/cluster/release_notes`
 - PowerHA SystemMirror for Smart Assists: `/usr/es/sbin/cluster/release_notes_assist`

PowerHA SystemMirror for IBM Systems Director

PowerHA SystemMirror is a plug-in for IBM Systems Director that uses a web browser interface to manage PowerHA SystemMirror capabilities.

PowerHA SystemMirror for IBM Systems Director, has the following functions:

- Discover existing PowerHA SystemMirror clusters.
- Collect information and a variety of reports about the state and configuration of applications and clusters.
- Receive live and dynamic status updates for clusters, sites, nodes, and resource groups.
- Single sign-on capability that you can use to have full access to all clusters with only one user ID and password.
- Access and search log files.
- Displays a summary page that you can view the overall status of all known clusters and resource groups.
- Create clusters and add resource groups with wizards.
- Apply updates to the PowerHA SystemMirror Agent using the Director Update Manager.

Related information:

 [IBM Systems Director Information Centers](#)

What's new in for PowerHA SystemMirror for IBM Systems Director

Read about new or significantly changed information for the PowerHA SystemMirror for IBM Systems Director topic collection.

How to see what's new or changed

In this PDF file, you might see revision bars (|) in the left margin that identify new and changed information.

December 2013

The following information is a summary of the updates made to this topic collection:

- You can use offline simulation mode and planning mode to explore the interface in a safe environment that does not affect your current cluster configuration. For more information, see the following topics:
 - “PowerHA SystemMirror for IBM Systems Director cluster simulator mode” on page 9
 - “Using offline simulation mode” on page 9
 - “Using planning mode” on page 10
- Added information about specific versions of that must be installed to use specific versions of PowerHA SystemMirror for IBM Systems Director. For more information, see “Prerequisites for installing PowerHA SystemMirror for IBM Systems Director” on page 5.
- You can restore a clusters configuration by using a snapshot. For more information, see “Restoring a cluster by using snapshot” on page 35.
- You can manually respond to a cluster split or merge. For more information, see “Manually responding to a cluster that uses a split policy or a merge policy” on page 38.

March 2013

The following information is a summary of the updates made to this topic collection:

- Added the following new commands:
 - “importcluster command” on page 49
 - “lsclrole command” on page 53
 - “lsefsKeystore command” on page 57
 - “lsfallbacktimer command” on page 58
 - “lsgroup command” on page 60
 - “lsldapcl command” on page 63
 - “lsldapsrv command” on page 64
 - “lssite command” on page 73
 - “lsversion command” on page 77
 - “mkdependency command” on page 87
 - “mkefsKeystore command” on page 90
 - “mkfallbacktimer command” on page 91
 - “mkgroup command” on page 95
 - “mkinterface command” on page 97
 - “mkldapcl command” on page 98
 - “mkldapsrv command” on page 100
 - “mkmethod command” on page 107
 - “mksecurity command” on page 113
 - “mkserviceip command” on page 115
 - “mktape command” on page 119
 - “mkuser command” on page 121
 - “moveresgrp command” on page 102
 - “rmdependency command” on page 133
 - “rmefsKeystore command” on page 135
 - “rmfallbacktimer command” on page 136
 - “rmgroup command” on page 139
 - “rmldapcl command” on page 141
 - “rmldapsrv command” on page 142
 - “rmmethod command” on page 145
 - “rmpersistentip command” on page 148
 - “rmserviceip command” on page 151
 - “rmsnapshot command” on page 154
 - “rmtape command” on page 155
 - “rmuser command” on page 156
 - “rssnapshot command” on page 158
 - “startcluster command” on page 159
 - “startnode command” on page 161
 - “stopcluster command” on page 164
 - “stopnode command” on page 166
 - “syncfilecollection command” on page 170
 - “vlog command” on page 174

- "lssite command" on page 73
- "mksite command" on page 116
- "rmsite command" on page 152
- "mkresgrp command" on page 109
- "lsvolume group command" on page 78

November 2012

The following information is a summary of the updates made to this topic collection:

- You can perform the following new tasks from the **Federated security** tab:
 - "Changing a user's password by using PowerHA SystemMirror for IBM Systems Director" on page 19
 - "Creating RBAC roles by using PowerHA SystemMirror for IBM Systems Director" on page 19
 - "Changing RBAC roles by using PowerHA SystemMirror for IBM Systems Director" on page 19
- You can perform the following new tasks from the **Reports** tab:
 - "Creating reports by using PowerHA SystemMirror for IBM Systems Director" on page 24
 - "Configuring report settings by using PowerHA SystemMirror for IBM Systems Director" on page 24
 - "Uploading reports by using PowerHA SystemMirror for IBM Systems Director" on page 25
 - "Emailing a report by using PowerHA SystemMirror for IBM Systems Director" on page 25
- You can perform the following new tasks from the **Storage** tab:
 - "Creating a Logical Volume Manager mirror pool" on page 33
 - "Adding a disk to an LVM mirror pool" on page 34
- You can perform the following new tasks from the **General** tab:
 - "Changing the name of a site" on page 41
 - "Bringing site services online or offline" on page 42
 - "Adding nodes to a site" on page 42
- You can perform the following new tasks from the **Events and Alerts** tab:
 - "Viewing event methods" on page 42
 - "Changing event method properties" on page 43
 - "Adding post-event and pre-event commands to event methods" on page 43
 - "Creating an event method summary report" on page 44
 - "Viewing event commands" on page 44
 - "Creating an event command" on page 44
 - "Deleting an event command" on page 45
 - "Viewing events in an event log by using PowerHA SystemMirror for IBM Systems Director" on page 45
- You can perform the following new tasks from the **HyperSwap and Replicated Storage** tab:
 - "Creating a replicated mirror group by using PowerHA SystemMirror for IBM Systems Director" on page 46
 - "Changing a replicated mirror group by using PowerHA SystemMirror for IBM Systems Director" on page 47
 - "Creating a storage agent for a cluster by using PowerHA SystemMirror for IBM Systems Director" on page 47
 - "Changing a storage agent by using PowerHA SystemMirror for IBM Systems Director" on page 48
 - "Creating a storage system by using PowerHA SystemMirror for IBM Systems Director" on page 48
 - "Changing a storage system by using PowerHA SystemMirror for IBM Systems Director" on page 48

- You can add Capacity on Demand (CoD) functions to an application controller. For more information about CoD functions, see the following topics:
 - “Adding Capacity on Demand for an application controller by using PowerHA SystemMirror for IBM Systems Director” on page 30
 - “Changing properties of Capacity on Demand for an application controller by using PowerHA SystemMirror for IBM Systems Director” on page 30
- You can perform the following new tasks from the **Other** tab:
 - “Viewing custom verification methods by using PowerHA SystemMirror for IBM Systems Director” on page 45
 - “Creating a custom verification method by using PowerHA SystemMirror for IBM Systems Director” on page 46
 - “Changing a custom verification method by using PowerHA SystemMirror for IBM Systems Director” on page 46

February 2012

The following information is a summary of the updates made to this topic collection:

- You can replace a repository disk with a new repository disk by using PowerHA SystemMirror for IBM Systems Director. For more information about replacing a repository disk, see “Viewing a repository disk with PowerHA SystemMirror for IBM Systems Director” on page 40.

December 2011

The following information is a summary of the updates made to this topic collection:

- Added information about configuring and managing resources for clusters and resource groups. For example, you can add application controllers, application monitors, service IP labels, and IP addresses to a resource group. For more information, see “Resources” on page 26.
- Added information about configuring and managing storage devices. For example, you can create a tape device for a cluster, add a file system to a resource group, and create a volume group. For more information see, “Storage” on page 31.
- Added information about federated security functionality. With federated security you can create a security group, security user, and a Lightweight Directory Access Protocol (LDAP) connection. For more information, see “Federated security” on page 15.
- Added information about a wizard that you can use to set up a SAP liveCache Hot Standby solution for your environment. For more information see, “Setting up SAP liveCache Hot Standby” on page 11.
- Added information about the following commands:
 - “lsresgrp command” on page 71
 - “mkresgrp command” on page 109
 - “mksnapshot command” on page 118
 - “rmresgrp command” on page 149

Installing PowerHA SystemMirror for IBM Systems Director

PowerHA SystemMirror Version 7.1 automatically recognizes previously discovered systems that are capable of using PowerHA SystemMirror, as long as they meet prerequisites.

Note: You can download different versions of the PowerHA SystemMirror for IBM Systems Director plug-in for free from the PowerHA SystemMirror for IBM Systems Director section on the IBM Systems Director Downloads - Plug-ins website. Before you install a new version of the PowerHA SystemMirror for IBM Systems Director plug-in, you must verify that your environment meets all prerequisite requirements.

Prerequisites for installing PowerHA SystemMirror for IBM Systems Director

Before you can install PowerHA SystemMirror for IBM Systems Director your environment must meet certain requirements.

The following table identifies the correct version of the AIX operating system, PowerHA SystemMirror, and IBM Systems Director that must be installed to use a specific version of PowerHA SystemMirror for IBM Systems Director.

Table 1. PowerHA SystemMirror for IBM Systems Director software requirements

PowerHA SystemMirror for IBM Systems Director version	PowerHA SystemMirror version	AIX version	IBM Systems Director version
7.1.0	7.1.0 or later	<ul style="list-style-type: none">• IBM AIX 6 with Technology Level 6, or later• IBM AIX Version 7.1, or later	6.2.0 or 6.2.1
7.1.1	7.1.1 or later	<ul style="list-style-type: none">• IBM AIX 6 with Technology Level 7, or later• IBM AIX 7 with Technology Level 1, or later	6.2.1
7.1.1.1	7.1.1 or later	<ul style="list-style-type: none">• IBM AIX 6 with Technology Level 7, or later• IBM AIX 7 with Technology Level 1, or later	6.3.0
7.1.2	7.1.2 or later	<ul style="list-style-type: none">• IBM AIX 6 with Technology Level 8, or later• IBM AIX 7 with Technology Level 2, or later	6.3.0 or later
7.1.3	7.1.3 or later	<ul style="list-style-type: none">• IBM AIX 6 with Technology Level 9, or later• IBM AIX 7 with Technology Level 3, or later	6.3.2 or later

Installing PowerHA SystemMirror for IBM Systems Director on the server

Before you install the PowerHA SystemMirror for IBM Systems Director plug-in, you must have installed the correct version of IBM Systems Director and PowerHA SystemMirror.

Installing the PowerHA SystemMirror for IBM Systems Director plug-in on AIX

To install the PowerHA SystemMirror for IBM Systems Director plug-in for AIX, complete the following steps:

1. Insert the installation media into the CD-ROM or DVD drive.
2. Log on to AIX system as root and from the command line run the following command:
`cd /`
3. Create a new directory by running the following command:
`mkdir /media`
4. Mount the CD-ROM or DVD device by running the following command:
`mount /dev/cd0 /media`
5. Go to the AIX directory on the media by running the following command:
`cd /media/director/AIX`

6. Start the PowerHA SystemMirror for IBM Systems Director plug-in installation wizard by running the following command:

```
./IBMSystemsDirector_PowerHA_sysmirror_Setup.bin
```

7. Follow the wizard instructions and prompts to complete the installation.

Installing the PowerHA SystemMirror for IBM Systems Director plug-in for Linux

To install the PowerHA SystemMirror for IBM Systems Director plug-in for the Linux operating system, complete the following steps:

1. Insert the installation media into the CD-ROM or DVD drive.
2. Log on to the Linux system as root and from the command line run the following command:

```
cd /
```
3. Create a new directory by running the following command:

```
mkdir /media
```
4. Mount the CD-ROM or DVD device by running the following command:

```
mount /dev/cd0 /media
```
5. Go to the Linux directory on the media by running the following command:

```
cd /media/director/Linux
```
6. Start the PowerHA SystemMirror for IBM Systems Director plug-in installation wizard by running the following command:

```
./IBMSystemsDirector_PowerHA_sysmirror_Setup.bin
```
7. Follow the wizard instructions and prompts to complete the installation.

Installing the PowerHA SystemMirror for IBM Systems Director plug-in for Windows

To install the PowerHA SystemMirror for IBM Systems Director plug-in for Windows, complete the following steps:

1. Log on to the system with either local or domain administrator authority.
2. Insert the installation media into the CD-ROM or DVD drive.
3. Browse the installation media and go to the `media/director/windows` directory.
4. Double-click the `IBMSystemsDirector_PowerHA_sysmirror_Setup.exe` file.
5. Follow the wizard instructions and prompts to complete the installation.

Related information:

 [IBM Systems Director Downloads - Plug-ins](#)

Installing PowerHA SystemMirror for IBM Systems Director on nodes

To install the PowerHA SystemMirror for IBM Systems Director file set, the system must be running IBM AIX 7 with Technology Level 1, or later, or IBM AIX 6 with Technology Level 7, or later.

To install the PowerHA SystemMirror for IBM Systems Director file set, complete the following steps:

1. Insert the installation media into the CD-ROM or DVD drive.
2. Log in to the AIX system as root and enter the following command from the command line: `smi t install_latest`
3. Specify `/dev/cd0` as the installation input device.
4. Specify `cluster.es.director.agent.rte` as the software to install.
5. Press Enter to install the PowerHA SystemMirror for IBM Systems Director client file set.

Unlocking nodes in IBM Systems Director after installation has completed

Unlocking a node automatically discovers, but does not request access to, all other nodes in the cluster. Those nodes that are part of a cluster that IBM Systems Director does not have access to will only be able to perform limited actions.

You need to unlock a node for one of the following reasons:

- To discover an existing PowerHA SystemMirror cluster. Unlocking the node triggers IBM Systems Director to automatically pull in the defined cluster and resource group definitions during discovery. Typically, these end points are in clusters that have been defined using the SMIT interface.
- To make an end point available for configuration as a cluster by using IBM Systems Director.
- To unlock an end point that is already part of a cluster to ensure full management capabilities by using IBM Systems Director.

Note: You should unlock all nodes in a cluster to take full advantage of the management capabilities.

After you have installed the PowerHA SystemMirror for IBM Systems Director plug-in, complete the following steps to unlock nodes:

Note: The following steps assume that you have already discovered the end points, but have not unlocked them with IBM Systems Director.

1. Login to your IBM Systems Director system as root.
2. Select **Navigate Resources** from the Navigation tree.
3. Select **All Operating Systems** from the Navigate Resources panel.
4. Right-click on each locked node and select **Security > Request Access**.
5. Enter the root username and password.

Logging in to the IBM Systems Director Console

You can access the PowerHA SystemMirror application from the IBM Systems Director console that is a graphical interface or from a command line interface.

To access PowerHA SystemMirror from the IBM Systems Director console, complete the following steps:

1. To start the console, open a browser and go to `https://HostName:8422/ibm/console` for secure access, or go to `http://HostName:8421/ibm/console` for normal access. **HostName** is your server name.
2. On the login page, enter your user name and password for the IBM Systems Director server.

Navigating the PowerHA SystemMirror for IBM Systems Director interface

With the PowerHA SystemMirror interface you can complete many different tasks and it provides unique views to help you manage your environment.

The following figure identifies the different areas of the graphical interface that are referenced to in this topic.

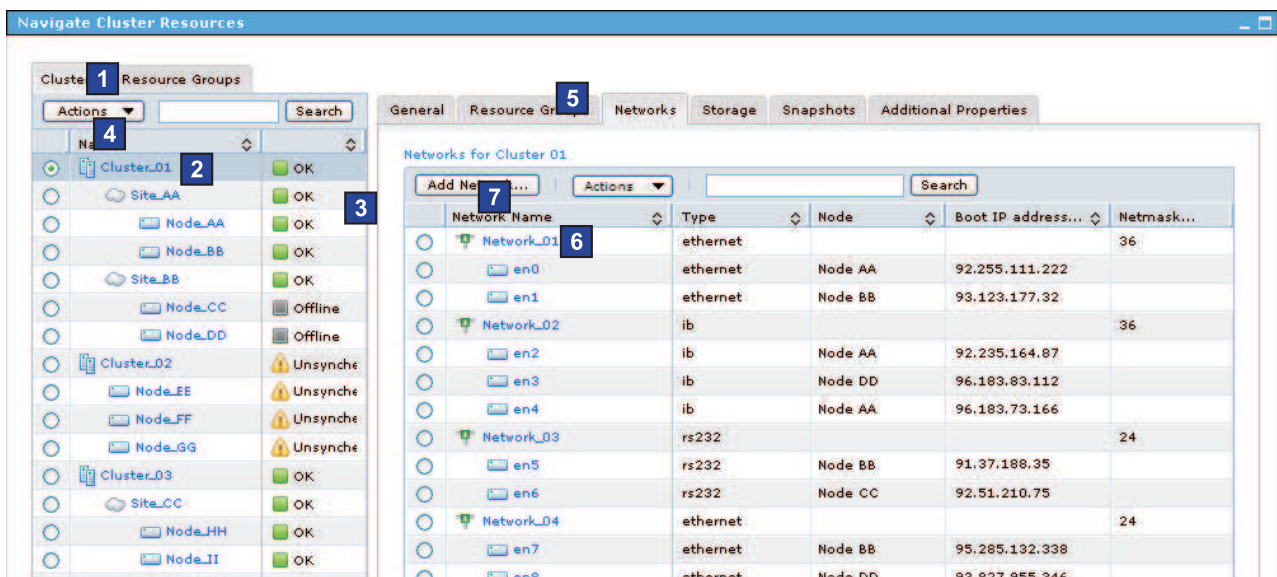


Figure 1. PowerHA SystemMirror graphical interface

1 Navigation area tabs

Clusters Tab: This tab provides a hierarchical view of each cluster known by IBM Systems Director based on the topology of each cluster.

Resource Groups Tab: This tab provides a similar hierarchical view, but its view is based on each cluster's known resource groups.

2 Navigation tree

This area uses indentation to provide a hierarchical view of the cluster topology or resource groups depending on what tab you select for **1**. You can left-click a specific item to obtain more information about it in the Content Area (**6**). To perform an action on a particular item, you can right-click the item to display a menu of actions currently available for that item.

3 Navigation tree status

The status area provides an at-a-glance indication of the health of a particular item, allowing you to easily identify problem areas.

4 Navigation area Actions button

With the **Actions** button you can perform an action on the item that is currently selected in the navigation tree (**2**). Click it to display a menu of actions currently available for that item. You can also enter text to search the navigation tree (**2**).

5 Content area tabs

These tabs provide you with views into the various functional areas of PowerHA SystemMirror. These tabs provide you with an uncluttered unified view into one particular aspect of the cluster.

6 Content area

This area provides a view of one particular functional area for the selected cluster or resource group. You can examine the specific details of your cluster and filter out everything else. You can take actions on many of the items displayed in this view by right-clicking the item.

7 Content area Actions button

With the **Actions** button you can perform actions on the item currently selected in the Content area (**6**). Clicking the **Actions** button creates a menu displaying the actions that are currently valid for the selected item.

Related concepts:

“Configuring clusters and resource groups by using PowerHA SystemMirror for IBM Systems Director” on page 11

Use PowerHA SystemMirror for IBM Systems Director to complete System Management Interface Tool (SMIT) tasks, such as creating clusters or resource groups, verifying and synchronizing a cluster, and moving a resource group to a different node.

PowerHA SystemMirror for IBM Systems Director cluster simulator mode

The cluster simulator mode is available in PowerHA SystemMirror for IBM Systems Director 7.1.3, or later. You can use offline simulation mode and planning mode to explore the interface in a safe environment that does not affect your current cluster configuration settings.

The cluster simulator function supports the following modes:

Online

In online mode, you can interact with the actual PowerHA SystemMirror cluster that is configured on agent nodes. When the system is running in the online mode, the cluster definition can be exported into an XML file.

Offline simulation

In offline simulation mode, changes that are made to the PowerHA SystemMirror cluster configurations are saved in an XML file instead of running it on the actual cluster of the agent nodes. This mode provides a safe sandbox to create clusters, to remove clusters, and to change other cluster configuration settings.

Planning

Planning mode is identical to offline simulation mode, except that the configuration data is collected from the actual cluster environment. In planning mode, the data about host names, disks, volume groups, and IP addresses are collected from the current cluster configuration settings. Because the cluster configuration uses real data that is collected from the environment, you can choose to deploy the cluster on the actual nodes.

The XML configuration files are in the `/opt/ibm/director/PowerHASystemMirror/eclipse/plugins/com.ibm.director.power.ha.systemmirror.common_7.1.3.0/bin/xml` directory on the IBM Systems Director server. Sample XML files are in the `/opt/ibm/director/PowerHASystemMirror/eclipse/plugins/com.ibm.director.power.ha.systemmirror.common_7.1.3.0/bin/xml/simu` directory.

Using offline simulation mode

You can use offline simulation mode to explore the PowerHA SystemMirror for IBM Systems Director interface and to save your PowerHA SystemMirror cluster configuration settings in an XML file.

In offline simulation mode, you can access all the interface elements and functions for PowerHA SystemMirror for IBM Systems Director.

To use offline simulation mode, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. From the **Mode** list, select **Offline Simulation**.

4. Select one of the following options:

Start a configuration from a simulation file

Select this option to use a sample XML file that is shipped with the product.

You can select one of the following sample configuration files:

NSC_7122.xml

Contains a cluster without a site that has two PowerHA SystemMirror 7.1.2.2 nodes.

NSC_713.xml

Contains a cluster without a site that has two PowerHA SystemMirror 7.1.3 nodes.

SC_713.xml

Contains a stretched cluster with two sites and four PowerHA SystemMirror 7.1.3 nodes.

LC_713.xml

Contains a linked cluster with two sites and four PowerHA SystemMirror 7.1.3 nodes.

Continue a configuration from a custom simulation file

Select this option if you already have your own custom configuration file.

5. Click **Ok**.

Note: A banner is displayed that identifies PowerHA SystemMirror for IBM Systems Director is in simulation mode, and specifies the XML file location where any cluster configuration changes are saved.

Using planning mode

Planning mode is identical to offline simulation mode, except that the configuration data is collected from your current cluster environment.

You can use planning mode to explore the interface and save your PowerHA SystemMirror cluster configuration settings in an XML file.

In planning mode, the data about host names, disks, volume groups, and IP addresses are collected from the current cluster configuration. Because the cluster configuration uses real data that is collected from the environment, you can deploy the cluster on the actual nodes.

For example, you can work in planning mode to create your cluster configuration and save it as an XML. You can test the configuration and verify whether it is working correctly. You can also use the saved planning XML file to deploy the cluster configuration in your environment.

To use planning mode, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. From the **Mode** list, select **Planning**.
4. Select one of the following options:

Start a configuration and save changes into a planning file

Select this option to create an XML file.

Continue a configuration from a custom simulation file

Select this option if you want to change an existing XML file that you already created.

5. Click **Ok**.

Note: A banner is displayed that identifies PowerHA SystemMirror for IBM Systems Director is in simulation mode, and specifies the XML file location where any cluster configuration changes are saved.

Configuring clusters and resource groups by using PowerHA SystemMirror for IBM Systems Director

Use PowerHA SystemMirror for IBM Systems Director to complete System Management Interface Tool (SMIT) tasks, such as creating clusters or resource groups, verifying and synchronizing a cluster, and moving a resource group to a different node.

Related concepts:

“Navigating the PowerHA SystemMirror for IBM Systems Director interface” on page 7

With the PowerHA SystemMirror interface you can complete many different tasks and it provides unique views to help you manage your environment.

Configuring clusters using PowerHA SystemMirror for IBM Systems Director

Using PowerHA SystemMirror for IBM Systems Director you can add a network to a cluster, create a snapshot, and create security user for the cluster.

Related concepts:

“Managing clusters by using PowerHA SystemMirror for IBM Systems Director” on page 36

You can use IBM Systems Director to manage various aspects of a cluster, such as modifying cluster properties, bringing a cluster offline or online, and collecting problem determination data.

Creating a cluster

Use the Create a Cluster wizard to help you create a highly available cluster for your applications. A highly available cluster lets you set up resources and applications in your system to remain available in case of an unexpected hardware malfunction or unresolved error, as a result your applications remain up and running.

To create a cluster that has sites, you must use PowerHA SystemMirror for IBM Systems Director 7.1.2, or later.

To create a new cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Create a Cluster**.
4. Follow the wizard instructions to complete the appropriate fields, and click **Finish**.

Note: If you did not select **Start Cluster Services** from the **Summary** page, you can right-click the cluster, and select **Bring Services Online** to start the cluster later.

Related concepts:

“Configuring resource groups using PowerHA SystemMirror for IBM Systems Director” on page 25

Using PowerHA SystemMirror for IBM Systems Director you can create a resource group, add a volume group to a resource group, and create an application controller.

Setting up SAP liveCache Hot Standby

Use the Setup SAP liveCache Hot Standby Solution wizard to configure the primary node and the takeover node to use the SAP software. PowerHA SystemMirror manages the SAP software.

Important: Review the information in the Prerequisites for using the PowerHA SystemMirrorSAP liveCache Hot Standby wizard topic in the AIX information center before you use the wizard.

To set up SAP liveCache Hot Standby for PowerHA SystemMirror, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Setup SAP liveCache Standby Solution**.
4. Complete the wizard as directed, and click **Finish**.
5. Verify and synchronize the cluster.

Related information:

Smart Assist for SAP liveCache Hot Standby

SAP liveCache Hot Standby with PowerHA SystemMirror

Networks

Networks provide access to the highly available applications of a cluster, detect and diagnose failures, and communicate with the daemons that are running on other nodes in the cluster. Interfaces are the physical connection of a node to a network.

Adding a network to a cluster:

Networks in PowerHA SystemMirror provide you with a way to logically group network interface cards and a way to restrict the placement of service addresses and persistent addresses. For example, networks can be used to make sure an application only uses adapters A, B, and C, and not use adapters X, Y, and Z.

To add a network to a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster to which you want to add a network.
5. Click the **Networks** tab in the content area.
6. Click the **Add Network** button.
7. Complete the fields and click **OK**.

Adding an interface to a network:

Interfaces can be added to networks to increase the capacity of the network.

To add an interface to a network, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. Select the cluster for which you want to add an interface.
4. Click the **Networks** tab in the content area.
5. Right-click the appropriate network in the table, and select **Add Interfaces**.
6. Complete all fields, and click **OK**.

Moving an interface to a different network:

Moving interfaces from one network to another will not be supported in this release. Users will have to remove an interface from one network, then add it to another network.

To move an interface to a different network, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.

3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to move an interface to a different network.
5. Click the **Networks** tab in the content area.
6. Right-click the interface for which you want to move to a different network, and select **Move Interface to Another Network**.
7. Select the network for which you want to move the interface to, and click **OK**.

Moving IP addresses between interfaces:

Moving service or persistent IP addresses between interfaces should be done when an adapter is experiencing problems and needs to be replaced, or when the adapter is simply being replaced to upgrade it.

To move an IP address between an interface, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. Select the cluster you want to move the IP address to another interface.
4. Click the **Networks** tab in the content area.
5. Right-click the appropriate interface, and select **Move IP Addresses Between Interfaces**.
6. Complete the fields and click **OK**.

Snapshot

A *snapshot* is a record of all the data that defines a particular cluster configuration from a specific moment in time. You can restore snapshots if you have problems with an existing cluster configuration and want to go back to a previous clusters configuration.

Creating a snapshot:

Creating a snapshot saves a current copy of the cluster configuration. This saved configuration can later be used to restore the current cluster configuration, if necessary. You can apply a cluster snapshot to either an active or inactive cluster.

When you create a snapshot, you are saving a copy of the cluster configuration as it currently exists on the targeted node. If the cluster has been synchronized, the snapshot represents the current cluster configuration on all nodes. If the cluster has not been synchronized, the snapshot represents only the local configuration on the controlling node. In the latter case, if that snapshot were restored, it would result in that local configuration being synchronized across the entire cluster, whether or not that was what was intended. Therefore, before you capture a snapshot, you should synchronize the cluster.

To create a cluster snapshot, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to create a snapshot of its configuration.
5. Click the **Snapshots** tab in the content area.
6. Click **Create a Snapshot** in the table header.
7. Complete all required fields, and click **OK**.

Related tasks:

“Restoring a snapshot” on page 14

A snapshot can be restored to return the cluster to the configuration of the saved snapshot.

Viewing snapshot data:

Before you restore a snapshot, you can view the details for the snapshots that you have already captured.

To view snapshot data, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the cluster for which want to view snapshot data.
5. Click the **Snapshots** tab in the content area.
6. Click the appropriate snapshot entry in the table. The properties page for the snapshot is shown with a breadcrumb trail leading back to the snapshot table.
7. Click the **Details** tab to view snapshot.

Restoring a snapshot:

A snapshot can be restored to return the cluster to the configuration of the saved snapshot.

The snapshot is restored across the entire cluster, and not just on the controlling node. The data is restored on the controlling node, and then synchronization is done immediately from that node to all the other nodes in the cluster.

If the snapshot you are using to restore the cluster configuration was not synchronized before you saved the snapshot, you will be propagating the configuration settings that were previously established only on the controlling node to the entire cluster.

To restore a snapshot of a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster that has the snapshot you want to restore.
5. Click the **Snapshots** tab in the content area.
6. Right-click the snapshot from the table that you want to restore and select **Restore**.
7. Verify the settings and click **OK**.

Related tasks:

“Creating a snapshot” on page 13

Creating a snapshot saves a current copy of the cluster configuration. This saved configuration can later be used to restore the current cluster configuration, if necessary. You can apply a cluster snapshot to either an active or inactive cluster.

Creating a custom snapshot method:

You can create a custom snapshot method to collect additional information about the cluster, or any other aspect of a node, such as operating system details or local application information. This information is appended to the **.info** file of the snapshot.

To create a custom snapshot method, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to create a custom snapshot method.

5. Click the **Snapshots** tab in the content area.
6. Select **Custom Snapshot Methods** from the **View** field.
7. Click **Create** in the table header.
8. Complete all required fields and click **OK**.

Federated security

To implement federated security, you can use role-based access control (RBAC) and Encrypted Files System (EFS) with PowerHA SystemMirror by using an existing LDAP configuration as the centralized information base for clusters. If you do not have an existing LDAP configuration, you can store the relevant information in the local file system. Using RBAC and Encrypted Files System (EFS) creates a high level of security for your cluster.

Related information:

Planning for federated security

Installing federated security

Creating a security group in a cluster:

A security group is a collection of users who can share access authority to protected resources. You can create a security group for users who access the same applications or the same hardware, and perform similar tasks.

To create a security group, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you are creating a security group.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **User Groups**.
7. Click **Create User Group**.

Note: You can right-click the security group and select **Create like** to create a new security group based on the existing security group properties.

8. Complete all required fields and click **OK**.

Changing properties for a security group in a cluster:

You can change the fields such as, the **Group ID**, **Registry**, and **Administrative group** fields, which are used to create the security group.

To change the properties for a security group in a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to change the security group properties.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **User Groups**.
7. Right-click the security group for which you want to change the properties, and click **Properties**.
8. Click **Edit**.
9. Complete all required fields and click **OK**.

Adding a security user to a security group:

A security user must have the same registry type as that of the security groups of which it is a member. A security user can be a member of up to 64 security groups.

To add a security user to a security group, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to add a security user to a security group.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **User Groups**.
7. Right-click the security group for which you want to add a security user, and click **Add a User**.
8. Select the security user from the table that you want to add to the selected security group, and click **OK**.

Note: If a security user that you want to add to the security group is not listed, that security user has one or more of the following problems:

- The security user is already part of the group.
- The security user does not have the same registry type as the security group.
- The security user already belongs to 65 or more security groups.
- The security user was not created.

Related tasks:

“Creating a security user for the cluster”

If the keystore properties of a security user are different from the properties of a security group, the property definitions of the security user are used.

Creating a security user for the cluster:

If the keystore properties of a security user are different from the properties of a security group, the property definitions of the security user are used.

To create a security user for the cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to create a security user.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **User Groups**.
7. Click **Create User**.
8. Complete all required fields and click **OK**.

Related tasks:

“Adding a security user to a security group”

A security user must have the same registry type as that of the security groups of which it is a member. A security user can be a member of up to 64 security groups.

Creating a new security user based on an existing security user:

When you create a new security user based on an existing security user, you can change or keep some of the existing properties.

To create a new security user based on an existing security user, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to create a new security user that is based on an existing security user.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **User Groups**.
7. Right-click the security user from the table, and click **Create Like**.
8. Complete all required fields and click **OK**.

Changing properties of a security user in a cluster:

You can change the same fields that you used to create the security group, such as, the **User ID**, **Home directory**, and **Primary group**.

To change the properties of a security user in a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to change the properties for a security user.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, click **Users**.
7. Right-click the security user for which you want to change the properties, and click **Properties**.
8. Click **Edit**.
9. Complete all required fields and click **OK**.

Changing a security user password:

If you do not want to change the password yourself for a security user, you can use an option to make the user change their password the next time they sign into the system.

To change a security user password, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to change the security user password.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **Users**.
7. Right-click the security user for which you want to change the password, and click **Change Password**.
8. Complete all required fields and click **OK**.

Creating an LDAP connection:

You can use a Lightweight Directory Access Protocol (LDAP) connection to create a connection between the server and the client on a node in the cluster.

After you create the server and client connection, LDAP can access and update information in a directory. LDAP is used by the cluster to keep authentication, groups, and user information common across the cluster.

To create an LDAP connection, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to create an LDAP server connection.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **LDAP management**.
7. Click **Create LDAP**.
8. Complete all required fields and click **OK**.

Disconnecting an LDAP connection:

You cannot change an LDAP connection. If you want to change an existing LDAP connection, you must disconnect the LDAP connection and create a new LDAP connection.

To disconnect an LDAP connection, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to disconnect the LDAP connection.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **LDAP management**.
7. Click **Disconnect LDAP**.
8. Review the information, and click **OK**.

Setting up an encrypted file system configuration:

You can use an encrypted file system (EFS) configuration to maintain the details of users and groups in a cluster.

You can choose to store the EFS keystore in a shared file system or on an LDAP server. The EFS configuration must be completed before creating or adding EFS to the cluster.

To set up an EFS configuration, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to set up an EFS configuration.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **EFS management**.
7. Click **Setup EPS**.
8. Complete all required fields and click **OK**.

Changing an encrypted file system configuration:

You can change some of the fields that are used to set up the encrypted file system (EFS) configuration, such as, the **Volung group for keystore** and **Service IP label / addressHome direcotry** fields. However, you cannot change the EFS administrator password.

To change an EFS configuration, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to change an EFS configuration.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **EFS management**.
7. Click **Edit**.
8. Complete all required fields and click **OK**.

Changing a user's password by using PowerHA SystemMirror for IBM Systems Director:

You can enable the capability of users to change their own passwords by selecting **Allow to Change Password** from the **Federated Security** tab.

To enable a user to change their password, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster that contains the user ID for which you want to allow to change their password.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **User Groups**.
7. Right-click the security user from the table, and select **Allow to Change Password**.
8. Click **OK**.

Creating RBAC roles by using PowerHA SystemMirror for IBM Systems Director:

You can create a role-based access control (RBAC) for a user in a user group. You can create an RBAC role only if you have a LDAP configuration for the cluster.

To create RBAC roles for a user, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster that contains the user ID for which you want to create an RBAC role.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **Roles**
7. Click **Create a Role**.

Note: This button is available only if you have an LDAP configuration for the cluster.

8. Complete all required fields and click **OK**.

Changing RBAC roles by using PowerHA SystemMirror for IBM Systems Director:

If you already created an RBAC role for a user, you can change the properties of the RBAC role.

To change RBAC roles for a user, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.

4. Select the cluster that contains the user ID for which you want to change an RBAC role.
5. Click the **Federated Security** tab in the content area.
6. From the **View** field, select **Roles**
7. Click **Edit**.
8. Complete all required fields and click **OK**.

Alerts

Alerts are automated notifications that can be associated with specific cluster events, such as a `node_down` event. When the event is triggered, the alerts are automatically sent in an email to the defined list of recipients.

Viewing alerts:

You can view all the alerts that are set up for the cluster.

To view alerts, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Events and Alerts** tab in the content area.
6. Select **Alerts** from the **View** field.

Creating alerts:

You can create alerts to notify you by email when a specific event occurs on a cluster.

To create alerts, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Events and Alerts** tab in the content area.
6. Select **Alerts** from the **View** field.
7. Click **Create** from the table header.
8. Complete all required fields and click **OK**.

Related tasks:

“Testing an alert”

You should periodically verify that the alerts you have configured are working properly.

Testing an alert:

You should periodically verify that the alerts you have configured are working properly.

To test an alert, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Events and Alerts** tab in the content area.

6. Select **Alerts** from the **View** field.
7. Right-click the alert you want to test and select **Test**.

Related tasks:

“Creating alerts” on page 20

You can create alerts to notify you by email when a specific event occurs on a cluster.

Viewing alert event subscriptions:

You can view all the alert event subscriptions that are being used for the cluster.

To view an alert event subscriptions, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Events and Alerts** tab in the content area.
6. Select **Alerts** from the **View** field.
7. Click the appropriate alert entry in the table. The properties page for the alert is shown with a breadcrumb trail leading back to the alerts table.
8. Click the **Events** tab to view the events subscriptions.

File collections

A *file collection* is a PowerHA SystemMirror container that manages one or more files as a single unit. These containers can be used to easily synchronize files across every node in the cluster. File collections can also be configured to automatically synchronize, ensuring that all occurrences of their files will be identical on all the nodes in the cluster.

PowerHA SystemMirror creates the HACMP_Files file collection and the Configuration_Files file collection. You cannot add, remove, or change the properties for these two file collections.

Viewing file collections:

You can view a file collection to see what files it currently manages and its current synchronization settings.

To view file collections, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the cluster of the file collection
5. Click the **Other** tab in the content area.
6. Select **File Collections** from the **View** field.

Creating a file collection:

You can only create file collections on a fully configured cluster with working communications between the nodes.

To create a file collection, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.

3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the cluster for which you want to create a file collection for.
5. Click the **Other** tab in the content area.
6. Select **File Collections** from the **View** field.
7. Click **Create** from the table bar.
8. Complete all required fields and click **OK**.

Related tasks:

“Adding files to a file collection”

A file can only be associated with one file collection at a time. The only exception is when files are associated with the HACMP_Files collection.

“Synchronizing a file collection”

During file collection synchronization, the remote node makes a backup copy of the original file before PowerHA SystemMirror copies the file collection. The file collection is stored in the `/var/hacmp/filebackup/` directory.

Adding files to a file collection:

A file can only be associated with one file collection at a time. The only exception is when files are associated with the HACMP_Files collection.

To add files to a file collection, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to add files to a file collection.
5. Click the **Other** tab in the content area.
6. Select **File Collections** from the **View** field.
7. Right-click the file collection you want to add files to and select **Add Files**.
8. Complete all required fields and click **Add Files** to add multiple paths to the file collection.
9. Click **OK**.

Related tasks:

“Creating a file collection” on page 21

You can only create file collections on a fully configured cluster with working communications between the nodes.

Synchronizing a file collection:

During file collection synchronization, the remote node makes a backup copy of the original file before PowerHA SystemMirror copies the file collection. The file collection is stored in the `/var/hacmp/filebackup/` directory.

To synchronize a file collection, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to synchronize a file collection on.
5. Click the **Other** tab in the content area.
6. Select **File Collections** from the **View** field.
7. Right-click the file collection you want to synchronize and select **Synchronize Files**.

Related tasks:

“Creating a file collection” on page 21

You can only create file collections on a fully configured cluster with working communications between the nodes.

Logs

With PowerHA SystemMirror for IBM Systems Director you can view different log files and collect log data.

Viewing log files:

You can view a wide variety of log files. Each log file explains a particular functional aspect of the cluster.

To view logs, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Reports and Logs** tab in the content area.
6. Select **Logs files** from the **View** field.

Changing cluster-wide settings:

If you want to change the cluster-wide settings for all log files in the cluster, you can use the global settings function.

To change the cluster-wide log settings, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Reports and Logs** tab in the content area.
6. From the **View** field, select **Log files**.
7. Click **Cluster-wide Log Settings**.
8. Change the appropriate fields and click **OK**.

Viewing log data:

You can view the contents of any PowerHA SystemMirror log. This can be useful for monitoring the internal operations of the cluster, and identifying any problems you are having.

To view log data, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to view log data.
5. Click the **Reports and Logs** tab in the content area.
6. Select **Logs files** from the **View** field.
7. Click the appropriate log entry in the table. The properties page of the log is shown with a breadcrumb trail leading back to the log table.

8. Click the **Details** tab from the content area.
9. Filter log data by typing the appropriate words in the entry field for content that you want to keep in the data list, and then click **Filter**.

Collecting log data for problem determination:

Collecting log data for problem determination can result in multiple large files. A compressed archive file is created on the controlling node for each node in the cluster. Each archive file contains all the log files collected from a particular node.

To collect log data for problem determination, complete the following steps:

Note: If there is not enough disk space to complete this operation, the operation is cancelled.

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to collect cluster log data.
5. Click the **Reports and Logs** tab in the content area.
6. Select **Logs files** from the **View** field.
7. Click **Collect Log Files** from the table header.
8. Complete all required fields and click **OK**.

Note: The completion of this operation can take a long time to complete, especially if you have many nodes in the cluster.

Reports for PowerHA SystemMirror for IBM Systems Director

You can use PowerHA SystemMirror for IBM Systems Director to create a report, to send a report in an email, or to upload a report to an FTP server.

Creating reports by using PowerHA SystemMirror for IBM Systems Director:

PowerHA SystemMirror for IBM Systems Director has existing reports that are configured and ready for you to use. These existing reports include network status, application availability, and environment configuration.

To create a report, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to create a report.
5. Click the **Reports and Logs** tab in the content area.
6. From the **View** field, select **Reports**.
7. Click **Create and Save**.
8. Complete all required fields and click **OK**.

Configuring report settings by using PowerHA SystemMirror for IBM Systems Director:

You must configure the settings for reports to identify the number of reports you want to have available and the location on the cluster where reports are stored.

To configure the settings for reports, complete the following steps:

1. Log in to the IBM Systems Director Console.

2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to configure the report settings.
5. Click the **Reports and Logs** tab in the content area.
6. From the **View** field, select **Reports**.
7. Click **Report Settings**.
8. Complete all required fields and click **OK**.

Uploading reports by using PowerHA SystemMirror for IBM Systems Director:

To upload a report you must know the FTP location, host name, user ID, and corresponding password.

To upload a report to an FTP location, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to upload a report.
5. Click the **Reports and Logs** tab in the content area.
6. From the **View** field, select **Reports**.
7. Right-click the report that you want to upload, and select **Upload**.
8. Complete all required fields and click **OK**.

Emailing a report by using PowerHA SystemMirror for IBM Systems Director:

You must know the Simple Mail Transfer Protocol (SMTP) server name and port number of the cluster to email a report.

To email a report, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to email a report.
5. Click the **Reports and Logs** tab in the content area.
6. From the **View** field, select **Reports**.
7. Right-click the report that you want to upload, and select **Send**.
8. Complete all required fields and click **OK**.

Configuring resource groups using PowerHA SystemMirror for IBM Systems Director

Using PowerHA SystemMirror for IBM Systems Director you can create a resource group, add a volume group to a resource group, and create an application controller.

After resources are ready for use by PowerHA SystemMirror, they can be assigned to resource groups for easier management as single entities. For example, if an application depends on one or more volume groups and a service IP address, these resources can be placed in a single resource group and managed as one unit.

Related concepts:

“Managing resource groups by using PowerHA SystemMirror for IBM Systems Director” on page 40
Using PowerHA SystemMirror for IBM Systems Director you can change cluster-wide properties, bring

resource groups offline or online, and move a resource group to a different node.

Related tasks:

“Creating a cluster” on page 11

Use the Create a Cluster wizard to help you create a highly available cluster for your applications. A highly available cluster lets you set up resources and applications in your system to remain available in case of an unexpected hardware malfunction or unresolved error, as a result your applications remain up and running.

Creating a resource group for a cluster

You can use the Create a Resource group wizard to create highly available resource groups. Using a highly available resource group, you can select nodes and add resources such as application controllers, application monitors, and service IP addresses.

To create a cluster that has sites, you must use PowerHA SystemMirror for IBM Systems Director 7.1.2, or later.

To create a resource group for a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. From the Resource Group Management section, click **Create a Resource Group**.
4. Follow the wizard instructions to complete all required fields, and click **Finish**.

Resources

After you have defined a resource group, you can add resources to it. When you add or change resources in a resource group, PowerHA SystemMirror for IBM Systems Director displays only valid choices that are based on the resource group management policy.

Related information:

PowerHA SystemMirror resources and resource groups

Viewing resources in a cluster or resource group:

You can view all of the service IP labels and IP addresses, application controllers, and application monitors for resource groups that are defined in a cluster.

To view resources in a cluster or in a resource group, complete the following steps.

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster or resource group for which you want to view the resources.
5. Click the **Resources** tab in the content area.

Creating a service IP label or IP address for a cluster:

When you create a service IP label or IP address for a cluster, you do not have to immediately add it to the resource group.

You can use service IP labels and IP addresses for a single resource group or for an entire cluster.

To create a service IP label or IP address for a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.

4. Select the cluster for which you want to create a service IP label or IP address.
5. Click the **Resources** tab in the content area.
6. Click **Actions**, and then select **Create a Resource > Service IP Label/IP address**.
7. Complete all required fields and click **OK**.
8. Verify and synchronize the cluster.

Changing properties of a service IP label or an IP address:

You can use a service IP label and IP address for a resource group or a cluster.

To change the properties of a service IP label or an IP address, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Click the **Resources** tab in the content area.
5. Select the cluster or resource group for which you want to change the properties of the service IP label or IP address.
6. Right-click the appropriate service **IP label** or **IP address** from the table, and click **Properties**.
7. Click **Edit**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Adding a service IP label or IP address to a resource group:

Service IP labels and IP addresses are used by a single resource group. A resource group can have multiple service IP labels or IP addresses that are not identical.

To add a service IP label or IP address to a resource group, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the resource group for which you want to add a service IP label or IP address.
5. Click the **Resources** tab in the content area.
6. Click **Actions**, and select **Add a Resource > Service IP Label / IP Address**.
7. Complete all required fields and click **OK**.
8. Verify and synchronize the cluster.

Creating an application controller for a cluster:

You can use the same application controller for different resource groups inside a cluster. A resource group can use multiple application controllers.

To create an application controller for a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to create the application controller.
5. Click the **Resources** tab in the content area.
6. Click **Actions**, and select **Create a Resource > Application Controller**.

7. Complete all required fields and click **OK**.
8. Verify and synchronize the cluster.

Changing properties of an application controller:

Changes that you make to an application controller applies to the cluster and all resource groups inside the cluster that use the specified application controller.

To change the properties of an application controller, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Click the **Resources** tab in the content area.
5. Select the cluster or resource group for which you want to change the properties of the application controller.
6. Right-click the application controller from the table that you want to change, and select **Properties**.

Note: If the application controller that you want to change is not displayed, the application controller does not exist in the cluster or resource group.

7. Click **Edit**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Adding an application controller to a resource group:

You can add the same application controller to different resource groups. A resource group can have multiple application controllers.

To add an application controller to a resource group, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the resource group for which you want to add an application controller.
5. Click the **Resources** tab in the content area.
6. Click **Actions**, and select **Add a Resource > Application Controller**.
7. Complete all required fields and click **OK**.
8. Verify and synchronize the cluster.

Creating an application monitor for a cluster:

Before you can create an application monitor, you must verify that the cluster has an application controller. An application monitor must be used by an application controller. An application controller can have multiple application monitors. You can use an application monitor with one or more resource groups inside a cluster.

To create an application monitor for a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to create an application monitor.

5. Click the **Resources** tab in the content area.
6. Click **Actions**, and select **Create a Resource > Application Monitor**.
7. Complete all required fields and click **OK**.
8. Verify and synchronize the cluster.

Adding an application monitor to an application controller:

To use an application monitor you must add it to an application controller. Application monitors are linked to the same resource group to which the parent application controllers are associated with.

To add an application monitor to an application controller, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster or resource group for which you want to add an application monitor to an application controller.
5. Click the **Resources** tab in the content area.
6. Right-click the application controller for which you want to add an application monitor, and select **Add an Application Monitor**.

Note: If the application monitor that you want to add to an application controller is not displayed, the application monitor was not created in the cluster or resource group.

7. Select the monitor and click **OK**.
8. Verify and synchronize the cluster.

Changing properties of an application monitor:

Changes that you make to an application monitor applies to all resource groups, clusters, and application controllers that use the specified application monitor.

To change the properties of an application monitor, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster or resource that has the application monitor for which you want to change the properties.
5. Click the **Resources** tab in the content area.
6. Right-click the application monitor from the table for which you want to change the properties, and select **Properties**.
7. Click **Edit**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Suspend or resume all application monitors in a resource group:

You can suspend or resume all application monitors that belong to an application controller. If you suspend all application monitors, the corresponding application controller does not monitor the specified information. Therefore, the specified information is not discovered by PowerHA SystemMirror for IBM Systems Director.

You cannot suspend or resume a specific application monitor that belongs to an application controller. You must suspend or resume all application monitors that belong to an application controller.

To suspend or resume an application monitor, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the resource group for which you want to suspend or resume the application monitor.
5. Click the **Resources** tab in the content area.
6. Complete the following steps:
 - a. If you want to suspend all application monitors, right-click the application controller and select **Suspend Application Monitors**.
 - b. If you want to resume all application monitors, right-click the application controller and select **Resume Application Monitors**.
7. Review all information, and click **OK**.

Adding Capacity on Demand for an application controller by using PowerHA SystemMirror for IBM Systems Director:

You can configure resources that use dynamic LPAR (DLPAR) and Capacity on Demand (CoD) operations to make these resources available to an application when it falls over to standby LPAR nodes.

Before you can add CoD for an application controller, you must verify whether the cluster has an application controller.

To add CoD for an application controller, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to add a CoD for an application controller.
5. Click the **Resources** tab in the content area.
6. Right-click the application controller for which you want add CoD, and select **Add Capacity on Demand**.
7. Select **I agree**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Changing properties of Capacity on Demand for an application controller by using PowerHA SystemMirror for IBM Systems Director:

You can change the properties of Capacity on Demand (CoD) for an application controller, such as Desired memory, Minimum virtual/read CPUs, and Desired processor units.

To change CoD for an application controller, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to change the CoD properties for an application controller.
5. Click the **Resources** tab in the content area.
6. Right-click the CoD property that you want to change, and click **Properties**.

7. Click **Edit**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Storage

PowerHA SystemMirror provides cluster wide storage management capabilities that you can use to manage files systems, logical volumes, and volume groups across a cluster.

Some of the PowerHA SystemMirror configurations require shared storage. These PowerHA SystemMirror configurations have disk subsystems that support access from multiple nodes.

Creating a volume group:

After you create a volume group for a cluster, you can add it to a resource group. A volume group can be managed by one resource group. A resource group can manage multiple volume groups.

To create a volume group, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you are creating a volume group.
5. Click the **Storage** tab in the content area.
6. From the **View** field, click **Volume groups and file systems**.
7. Click **Create a Volume Group** to launch the Create a Volume Group wizard.
8. Complete all required fields in the wizard, and click **Finish**.
9. Verify and synchronize the cluster.

Adding a volume group to a resource group:

You can add a volume group to any resource group in a cluster that are not being used by another resource group. Volume groups can be managed by a single resource group. A single resource group can manage multiple volume groups.

To add a volume group to a resource group, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the resource group for which you want to add a volume group.
5. Click the **Storage** tab in the content area.
6. From the **View** field, click **Volume Groups**.
7. Click **Add a volume group**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Adding a raw disk to a resource group:

You can add a raw disk to any resource group in a cluster that are not being used by another resource group. Raw disks can be used by a single resource group. A single resource group can have multiple disks.

To add a raw disk to a resource group, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the resource group for which you want to add a disk.
5. Click the **Storage** tab in the content area.
6. From the **View** field, select **Raw disks**.
7. Click **Add a disk**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Creating a tape device for a cluster:

After you create a tape device for a cluster, you can add it to a resource group. A tape device can only be used by one resource group.

To create a tape device, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to create a tape device.
5. Click the **Storage** tab in the content area.
6. From the **View** field, click **Tape**.
7. Click **Create a tape resource**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Related information:

Planning shared disk and tape devices

Adding a tape device to a resource group:

You can add a tape device to any resource group in a cluster that are not being used by another resource group. Tape devices can be used by a single resource group. A single resource group can have multiple tape devices.

To add a tape device to a resource group, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the resource group for which you want to add a tape device.
5. Click the **Storage** tab in the content area.
6. From the **View** field, select **Tape**.
7. Click **Add a tape resource**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Changing properties of a tape device:

After you have created a tape device, you can change the properties of the tape device at any time.

To change the properties of a tape device, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Click the **Storage** tab in the content area.
5. From the **View** field, click **Tape**.
6. Right-click the tape device from the table for which you want to change the properties, and click **Properties**.
7. Click **Edit**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Creating a file system for a cluster:

After you create a file system for a cluster, you can add it to a resource group. A file system can be used by one resource group.

To create a file system, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to create a file system.
5. Click the **Storage** tab in the content area.
6. From the **View** field, click **Volume groups and file systems**.
7. Right-click the volume group for which you want to create the file system, and click **Add a file system**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Creating a Logical Volume Manager mirror pool:

You can use mirror pools to group physical volumes into a scalable volume group. This process creates a mirror copy of a logical volume that is restricted to allocate physical volumes from a specified group.

To create a Logical Volume Manager (LVM) mirror pool, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to create an LVM mirror pool.
5. Click the **Storage** tab in the content area.
6. From the **View** field, select **Volume groups and LVM mirror pools**.
7. Click **Create a Mirror Pool**.

Adding a disk to an LVM mirror pool:

If disks are selected that do not already belong to the specified volume group, those disks are added to the volume group before joining the mirror group. If the disk that you want to add is not available, the disk is already associated with a mirror pool, a replicated mirror group, or a different volume group.

To add a disk to a Logical Volume Manager (LVM) mirror pool, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to add a disk to the LVM mirror pool.
5. Click the **Storage** tab in the content area.
6. From the **View** field, select **Volume groups and LVM mirror pools**.
7. Right-click the LVM mirror pool for which you want to add a disk, and select **Add a Disk**.
8. Select one or more disk from the table, and click **OK**.

Resource group dependencies

A variety of relationships, known as dependencies, can be established between resource groups to control where they run and how they start up. For example, a dependency can be established between two resource groups to ensure that those groups always run on the same node.

The startup order for multiple resource groups can be specified by using dependencies, such as ResourceGroup1 must always start before ResourceGroup2.

Dependencies that dictate where a set of resource groups run in relation to each other are known as *location dependencies*. Dependencies that dictate when a set of resources groups run in relation to each other are known as *temporal dependencies*.

Viewing resource group dependencies:

You can view all the resource group dependencies that are being used for the cluster.

To view resource group dependencies, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Dependencies** tab in the content area to display the resource group dependencies.

Creating a resource group dependency:

Consider creating a resource group dependency if the operation of one resource group is affected by the operation of another, or if one is dependent on the other.

To create a resource group dependency, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. From the Resource Group Management section, click **Manage Resource Groups**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Dependencies** tab in the content area to display the resource group dependencies.
6. Click **Actions** and select **Create Dependency**.

7. Select one of the following dependencies:

Note: Only one Online on Different Node dependency can exist per cluster. If such a dependency already exists for the cluster, the Online on Different Node option is not displayed.

Online on same site

In this type of dependency, a location policy is established so that the selected resource groups are started and remains online on the same site.

Parent-Child

In this type of dependency, during resource group acquisition, the parent resource group on a node is acquired before the children resource groups are acquired. During resource group release, children resource groups are released before the parent resource group is released. The parent resource group must be online on any node in the cluster before children (dependent) resource groups can be activated on a node.

Select

Online on Different Node

In this type of dependency, sets of resource groups are assigned High, Medium, and Low priority values, which are used in case there is contention for a given node at any point in time. Higher priority resource groups take precedence over lower priority groups at startup, fallover, and fallback. Each resource group in a set must have a different home node for startup. Only one Online On Different Nodes Dependency set is allowed per cluster.

Online on Same Node

In this type of dependency, a set of two or more resource groups are acquired and brought online on the same node according to the startup policy and the availability of the node required. On fallback and fallover, the set of resource groups are processed simultaneously and brought online on the same target node using the fallover and fallback policy defined for these groups.

Start After

In this type of dependency, during resource group acquisition, the target resource group on a node is acquired before the dependent resource group is acquired. The target resource group must be online on any node in the cluster before the dependent resource group can be activated on a node.

Stop After

In this type of dependency, during resource group release, the target resource group on a node is released before the dependent resource group is released.

8. Select the desired resource group dependencies and click **OK**.

Restoring a cluster by using snapshot

To restore the cluster configuration settings by using a snapshot, you must already have a snapshot that is saved from the version of PowerHA SystemMirror that is used on your system.

You must have installed PowerHA SystemMirror for IBM Systems Director Version 7.1.3, or later, to restore a snapshot.

Note: You cannot restore a PowerHA SystemMirror 6.1 snapshot by using PowerHA SystemMirror for IBM Systems Director.

To restore the cluster configuration settings by using a snapshot, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Restore Snapshot**.
4. Follow the wizard instructions to complete the appropriate fields, and click **Finish**.

Managing clusters and resource groups by using PowerHA SystemMirror for IBM Systems Director

Using PowerHA SystemMirror for IBM Systems Director you can change properties for clusters and resource groups.

Managing clusters by using PowerHA SystemMirror for IBM Systems Director

You can use IBM Systems Director to manage various aspects of a cluster, such as modifying cluster properties, bringing a cluster offline or online, and collecting problem determination data.

Related concepts:

“Configuring clusters using PowerHA SystemMirror for IBM Systems Director” on page 11

Using PowerHA SystemMirror for IBM Systems Director you can add a network to a cluster, create a snapshot, and create security user for the cluster.

Changing cluster properties

You can change the security settings, tuning options, and the split and merge policy properties for the cluster.

To change cluster properties, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the cluster for which you want to change the cluster properties.
5. From the **General** tab, and click **Edit**.
6. Complete the fields, and click **OK**.

Note: Any changes that you make must be propagated to all nodes by verifying and synchronizing the cluster.

Related tasks:

“Verifying and synchronizing a cluster”

When you make any changes to the cluster you should verify that the changes are working properly, and then you must synchronize the controlling node to propagate the changes to all the other nodes in the cluster.

Verifying and synchronizing a cluster

When you make any changes to the cluster you should verify that the changes are working properly, and then you must synchronize the controlling node to propagate the changes to all the other nodes in the cluster.

The verification process validates your changes, and the synchronization process ensures that all nodes within the cluster are updated with your changes. Synchronizing a cluster could delete or alter information in the cluster. You should save a snapshot of the cluster before you synchronize any changes.

To verify and synchronize a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Right-click the cluster you want to verify and synchronize, and select **Verify and Synchronize**.

Related tasks:

“Changing cluster properties” on page 36

You can change the security settings, tuning options, and the split and merge policy properties for the cluster.

“Undoing local changes of the configuration”

Undoing local changes of the configuration since the last synchronization saves a snapshot of the cluster before performing this operation. All user-configured PowerHA SystemMirror information since the last synchronization is deleted.

Undoing local changes of the configuration

Undoing local changes of the configuration since the last synchronization saves a snapshot of the cluster before performing this operation. All user-configured PowerHA SystemMirror information since the last synchronization is deleted.

To undo local changes of the configuration since the last synchronization, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Right-click the cluster where you want to discard all modifications made since the last synchronization, and select **Recovery > Undo Local Changes of the Configuration**.
5. Click **OK**. A message is displayed stating that any changes have been successfully undone since the last synchronization.

Related tasks:

“Verifying and synchronizing a cluster” on page 36

When you make any changes to the cluster you should verify that the changes are working properly, and then you must synchronize the controlling node to propagate the changes to all the other nodes in the cluster.

Bringing a cluster offline or online

When you bring cluster services online or offline, PowerHA SystemMirror activates or deactivates the resources according to how they were defined, and takes into consideration application dependencies, application start and stop scripts, dynamic attributes, and other parameters.

PowerHA SystemMirror uses the application monitors and application startup script to verify whether it needs to start the application or whether the application is already running.

To bring a cluster online or offline, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Right-click the cluster you want to bring online or offline.
5. Complete one of the following items:
 - a. If you want to bring services online, select **Bring Services Online**.
 - b. If you want to bring services offline, select **Bring Services Offline**.

Recovering from an event failure

Recovering from an event failure only advances the cluster manager to the next event following an error condition and does not correct any problems that led to the error. Therefore, it might be necessary to verify that all problems that led to the error condition have been addressed prior to running this action.

To recover from an event failure, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.

3. In the Cluster Management section, click **Manage Clusters**.
4. Right-click the cluster that you want to recover from an event failure, and select **Recovery > Recover from Event Failure**.
5. Read the warning message and if you want to continue, click **OK**.

Manually responding to a cluster that uses a split policy or a merge policy

If your site goes offline and you do not want the cluster to implement a split policy or a merge policy, you can manually specify the site that you want to be brought back online.

To manually respond to a cluster that goes offline and uses a split policy or a merge policy, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Right-click the cluster that you do not want to use a split policy or a merge policy, and select **Recovery > Manual response to cluster split or merge**.
5. Select the site that recovers the cluster and click **OK**.

Related information:

Split policy

Merge policy

Collecting problem determination data

Collecting problem determination data collects the cluster log file data that is necessary for problem determination and analysis. This action creates a very large log file; if there is not enough disk space to complete this operation, the operation will be cancelled.

To collect problem determination data, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Right-click the cluster for which you want to collect problem determination data, and select **Recovery > Collect Problem Determination Data**.
5. Verify debugging and log files settings, choose nodes to collect data from, and choose the node and directory to store the collected log file data, and then click **OK**.

Releasing cluster modification locks

Releasing cluster modification locks allows the cluster to be reconfigured in the event that the Staging Configuration Directory remains on a node after a failed dynamic reconfiguration.

You should only perform this task if you are having problems with your cluster. The most common problem occurs when you cannot synchronize your cluster. This problem can occur when a previous synchronization was interrupted by an unfortunate event, such as a node or network failure. Releasing the cluster locks can restore the ability to synchronize the cluster, but at a cost of possibly losing the changes that were not completely synchronized before. It is recommended that you only perform this action on the advice and guidance of IBM Service.

To release cluster modification locks, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.

4. Right-click the cluster for which you want to release the modification locks, and select **Recovery > Release Cluster Modification Locks**.
5. Click **OK** to release the cluster modification lock.

Refreshing cluster data

If the cluster data that you are viewing in PowerHA SystemMirror for IBM Systems Director is not displaying correct information about your environment, you can update the display by refreshing the cluster data.

To refresh cluster data, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Right-click the cluster that you want to refresh, and select **Refresh Cluster Data**.

Adding nodes to a cluster

You can add new nodes to a cluster after the cluster has been created without using the Create a Cluster wizard.

To add nodes to an existing cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Right-click the cluster that you want to add nodes to, and select **Add Nodes**.
5. Select one or more nodes from the table and click **OK**.

Adding a persistent IP label or a persistent IP address to a cluster

A persistent IP label or a persistent IP address is an IP alias that you can assign to a network for a specified node in a cluster.

Assigning a persistent IP label or IP address on a node allows you to have a node-bound address on a cluster that you can use to access a specific node in the cluster.

To add a persistent IP label or a persistent IP address to a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the node for which you are adding a persistent IP label or a persistent IP address.
5. Click the **Persistent IP** tab in the content area.
6. In the table, select the network connection for which you want to add a persistent IP label or a persistent IP address to a cluster.
7. Click **Actions**, and select **Add a Persistent IP**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Adding a repository disk to a cluster by using PowerHA SystemMirror for IBM Systems Director

You can add a repository disk to a cluster only if the cluster is running PowerHA SystemMirror 7.1.2, or later, Enterprise Edition for AIX.

To add a repository disk to a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Other** tab in the content area.
6. From the **View** field, select **Repository Disks**.
7. Click **Add backup repository disk**.
8. Complete all required fields and click **OK**.

Viewing a repository disk with PowerHA SystemMirror for IBM Systems Director

Cluster Aware AIX (CAA) detects when a repository disk failure occurs and sends a notification message. You will continue to receive notification messages until you replace the failed repository disk with a new repository disk.

The cluster operates in a restricted mode until you replace the failed repository disk. Therefore, do not attempt any cluster changes until you replace the failed repository disk.

To view a repository disk for a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click **Other** tab in the content area.
6. Select **Repository Disks** from the **View** field.

Related information:

Planning for repository disk

Repository disk failure

Managing resource groups by using PowerHA SystemMirror for IBM Systems Director

Using PowerHA SystemMirror for IBM Systems Director you can change cluster-wide properties, bring resource groups offline or online, and move a resource group to a different node.

Related concepts:

“Configuring resource groups using PowerHA SystemMirror for IBM Systems Director” on page 25

Using PowerHA SystemMirror for IBM Systems Director you can create a resource group, add a volume group to a resource group, and create an application controller.

Changing cluster-wide properties for a resource group

You can change cluster-wide properties, such as resource group settling time, resource group timeout, and resource group processing order.

To change cluster-wide properties, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to change the cluster-wide properties, and click **General**.
5. Click the **Edit** tab in the content area.
6. Complete all required fields and click **OK**.

Bringing resource groups online or offline

If you need to apply maintenance or service to an application, or other resources managed by a resource group, the resource group must be brought offline. Bringing the resource group offline prevents an unintended fallover (for example, the resource group might come online on another node) during the course of the maintenance. When you are ready for the resources within the resource group to be made available again, you must bring the resource group back online.

To bring resource groups online or offline, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. From the Resource Group Management section, click **Manage Resource Groups**.
4. Right-click the resource group that you want to bring online or offline.
5. Complete one of the following actions:
 - a. If you want to bring a resource group online, click **Bring Online**.
 - b. If you want to bring a resource group offline, click **Bring Offline**.
6. Complete all appropriate fields and click **OK**.

Moving a resource group to a different node

A resource group can be manually moved to a new node, just as it can be made to move automatically when a problem is detected and fallover occurs.

When a resource group is moved, all of the resources that it encompasses are moved with it, such as an application, any service IP addresses, volume groups, or file systems. Any applications that are moved are brought online on the target node, and any service IP addresses are aliased to the adapters on the target node.

To move a resource group to a different node, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. From the Resource Group Management section, click **Manage Resource Groups**.
4. Right-click the resource group that you want to move and select **Move to a Different Node**.
5. Complete the appropriate fields and click **OK**.

Managing sites by using PowerHA SystemMirror for IBM Systems Director

Your environment must be running PowerHA SystemMirror 7.1 Enterprise Edition for AIX to manage sites. By using PowerHA SystemMirror for IBM Systems Director, you can manage site properties, bring site services online or offline, and add a node to a site.

Changing the name of a site

You can only change the **Name** field of the site by using PowerHA SystemMirror for IBM Systems Director.

To change the name of the site, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Select the site for which you want to change the name.
5. From the **General** tab, click **Edit**.
6. Enter the new name for the site and click **OK**.

Bringing site services online or offline

When you bring site services online or offline, PowerHA SystemMirror activates or deactivates the resources according to how they are defined. PowerHA SystemMirror also considers application dependencies, application start and stop scripts, dynamic attributes, and other parameters when bringing site services online or offline.

To bring a site service online or offline, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Right-click the site that you want to bring online or offline.
5. Complete one of the following actions:
 - a. If you want to bring a site service online, click **Bring Cluster Services Online**.
 - b. If you want to bring a site service offline, click **Bring Cluster Services Offline**.
6. Complete all appropriate fields and click **OK**.

Adding nodes to a site

Your environment must be running must have PowerHA SystemMirror Version 7.1.2, or later, to define sites. However, you do not need this version to add nodes to an existing site.

A *site* is a set of nodes with its own policies, its own shared storage, and a cluster that can contain two sites. Each site is intended to be geographically separated from the other site. The sites are geographically separated to protect your applications from a disaster, such as a major power outage, a network outage, a flood, or fire. In a properly configured cluster, when such a disaster occurs, any resource groups being serviced by the disabled site are automatically brought up on the other site, which minimizes the effect of the disaster on your business operations. To increase the capacity of a site, you can add new nodes to the site.

To add nodes to an existing site, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. Right-click the site to which you want to add nodes, and select **Add Nodes**.
5. Select one or more nodes from the table and click **OK**.

Managing events by using PowerHA SystemMirror for IBM Systems Director

By using PowerHA SystemMirror for IBM Systems Director, you can manage event methods, event commands, and event logs.

Viewing event methods

When a cluster event occurs, the cluster manager runs the corresponding event method for that event. When the event method is being processed, a series of subevent scripts can run. PowerHA SystemMirror provides a method for each event and subevent.

To view event methods, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Events and Alerts** tab in the content area.

6. From the **View** field, select **Event Methods**.

Note: PowerHA SystemMirror must process all available data to view the event method. This process can take awhile to complete.

Changing event method properties

The cluster manager can recognize a specific series of events and subevents that creates a flexible customization scheme. If you customize event processing, you can provide the most efficient path to critical resources if a failure occurs.

To change an event method, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Events and Alerts** tab in the content area.
6. From the **View** field, select **Event Methods**.

Note: PowerHA SystemMirror must process all available data to view the event method. This process can take awhile to complete.

7. Click **Edit**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Adding post-event and pre-event commands to event methods

You can make event processing specific for your PowerHA SystemMirror environment by specifying event commands that run before and after a specific event is generated by the cluster manager.

You must create an event command before you can add a pre-event or post-event command.

To add post-event and pre-event commands to event methods, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Events and Alerts** tab in the content area.
6. From the **View** field, select **Event Methods**.

Note: PowerHA SystemMirror must process all available data to view the event method. This process can take awhile to complete.

7. Right-click an event method, and select **Properties**.
8. Click the **Pre- and Post-Event Commands** tab from the content area.
9. Complete one of the following actions:
 - a. If you want to add a pre-event command, click **Add Pre-event Command**.
 - b. If you want to add a post-event command, click **Add Post-event Command**.
10. Complete all required fields and click **OK**.
11. Verify and synchronize the cluster.

Related tasks:

“Creating an event command” on page 44

To create an event command, you must know the specific path to the script file that contains the

custom-defined event command. The script must be available on the node in the cluster.

Creating an event method summary report

After you create the event summary report, you can save it to view at a later time.

To create an event method summary report, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Events and Alerts** tab in the content area.
6. From the **View** field, select **Event Methods**.

Note: PowerHA SystemMirror must process all available data to view the event method. This process can take awhile to complete.

7. Click **Create Event Summary Report**.

Managing event commands

Event commands are user-defined commands that, when added to a cluster event, can run before or after the cluster event command runs.

Viewing event commands:

You can view all event commands that belong to a specific cluster by using PowerHA SystemMirror for IBM Systems Director.

To view event commands, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Events and Alerts** tab in the content area.
6. From the **View** field, select **Event Commands**.

Creating an event command:

To create an event command, you must know the specific path to the script file that contains the custom-defined event command. The script must be available on the node in the cluster.

To view event commands, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Events and Alerts** tab in the content area.
6. From the **View** field, select **Event Commands**.
7. Click **Create**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Related tasks:

“Adding post-event and pre-event commands to event methods” on page 43

You can make event processing specific for your PowerHA SystemMirror environment by specifying event commands that run before and after a specific event is generated by the cluster manager.

Deleting an event command:

You can delete an event command from PowerHA SystemMirror for IBM Systems Director. However, this process does not delete the script that the command was using from the node.

To delete an event command, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Events and Alerts** tab in the content area.
6. From the **View** field, select **Event Commands**.
7. Click **Delete**.

Viewing events in an event log by using PowerHA SystemMirror for IBM Systems Director

You can view events in an event log for a cluster by using PowerHA SystemMirror for IBM Systems Director.

To view an event in an event log, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Events and Alerts** tab in the content area.
6. From the **View** field, select **Event Log**.

Managing custom verification methods by using PowerHA SystemMirror for IBM Systems Director

You can use a custom verification method to check for a particular problem in your cluster.

Viewing custom verification methods by using PowerHA SystemMirror for IBM Systems Director

You can view custom verification methods for a cluster by using PowerHA SystemMirror for IBM Systems Director.

To view custom verification methods for a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Other** tab in the content area.
6. From the **View** field, select **Custom Verification Methods**.

Creating a custom verification method by using PowerHA SystemMirror for IBM Systems Director

You can create a custom verification method to troubleshoot problems by using PowerHA SystemMirror for IBM Systems Director.

To create a custom verification method for a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Other** tab in the content area.
6. From the **View** field, select **Custom Verification Methods**.
7. Click **Create**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Changing a custom verification method by using PowerHA SystemMirror for IBM Systems Director

You can change the **Name**, **Description**, and **Path** fields for a custom verification method.

To change a custom verification method for a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Cluster Management section, click **Manage Clusters**.
4. From the navigation tree, select the appropriate cluster.
5. Click the **Other** tab in the content area.
6. From the **View** field, select **Custom Verification Methods**.
7. Select a custom verification method from the table, and click **Edit**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Managing replicated storage

For resource groups in the cluster, you can use PowerHA SystemMirror for IBM Systems Director to manage replicated mirror groups, storage agents, and storage systems.

Creating a replicated mirror group by using PowerHA SystemMirror for IBM Systems Director

Replicated mirror groups provide an automated failover support, fallback support, and data replication for applications across sites in different geographic locations. You can use the Create a Replicated Mirror Group wizard to create a replicated mirror group.

Replicated mirror groups are supported only in PowerHA SystemMirror 7.1 Enterprise Edition for AIX.

You can use only the following storage technologies with replicated mirror groups:

- IBM DS8000® Metro Mirror (Inband)
- IBM DS8700 Global Mirror
- SAN Volume Controller (SVC)
- EMC SRDF
- XIV® Remote Mirror

- Hitachi Truecopy
- Hitachi Universal Replicator (HUR)

To create a replicated mirror group, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Create a Replicated Mirror Group** to launch the Create a Replicated Mirror Group wizard.
4. Follow the wizard instructions to complete all required fields, and click **Finish**.
5. Verify and synchronize the cluster.

Changing a replicated mirror group by using PowerHA SystemMirror for IBM Systems Director

You cannot change the Storage Type or Mirror Group Type fields for a replicated mirror group.

To change the properties for a replicated mirror group, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to change a replicated mirror group.
5. Click the **HyperSwap and Replicated Storage** tab in the content area.
6. From the **View** field, select **Replicated Mirror Groups**.
7. Right-click the mirror group, and select **Properties**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Creating a storage agent for a cluster by using PowerHA SystemMirror for IBM Systems Director

Storage agents represent a physical and logical component that is the interface for storage management-related APIs. These APIs are used by PowerHA SystemMirror to manage the replication for a disk. Typically, storage agents use a TCP/IP address as its method for communication.

You can use only the following storage technologies with storage agents:

- IBM DS8700 Global Mirror
- XIV Remote Mirror

To create a storage agent for a cluster, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to create a storage agent.
5. Click the **HyperSwap and Replicated Storage** tab in the content area.
6. From the **View** field, select **Storage Agents**.
7. Click **Create a storage agent**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Changing a storage agent by using PowerHA SystemMirror for IBM Systems Director

You cannot change the Storage Type field for a storage agent.

To change the properties for a storage agent, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to change a storage agent.
5. Click the **HyperSwap and Replicated Storage** tab in the content area.
6. From the **View** field, select **Storage Agents**.
7. Right-click the storage agent that you want to change, and click **Properties**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Creating a storage system by using PowerHA SystemMirror for IBM Systems Director

Storage systems represent the storage unit and controller that hosts or controls one or more disks or volumes groups that are being replicated. Typically, every storage system has a (manufacturer-specific) unique identifier. For example, IBM.2107-1300851 F is the unique identifier for an IBM DS8300 storage system.

You can use only the following storage technologies with replicated mirror groups:

- IBM DS8000 Metro Mirror (Inband)
- IBM DS8700 Global Mirror
- SAN Volume Controller (SVC)
- XIV Remote Mirror

To create a storage system, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to create a storage system.
5. Click the **HyperSwap and Replicated Storage** tab in the content area.
6. From the **View** field, select **Storage Systems**.
7. Click **Create a storage system**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

Changing a storage system by using PowerHA SystemMirror for IBM Systems Director

You can change the **Name** and **Storage Agents** fields for a storage system.

To change the properties of a storage system, complete the following steps:

1. Log in to the IBM Systems Director Console.
2. On the Welcome page, click the **Plug-ins** tab and click **PowerHA SystemMirror Management**.
3. In the Resource Group Management section, click **Manage Resource Groups**.
4. Select the cluster for which you want to change the properties of a storage system.
5. Click the **HyperSwap and Replicated Storage** tab in the content area.

6. From the **View** field, select **Storage Systems**.
7. Right-click the storage system for which you want to change the properties, and click **Properties**.
8. Complete all required fields and click **OK**.
9. Verify and synchronize the cluster.

PowerHA SystemMirror for IBM Systems Director commands

You can create scripts that use the PowerHA SystemMirror commands to manage your clusters.

You can use the following commands to manage your cluster.

help command

Purpose

Provides the ability to obtain help for PowerHA SystemMirror for IBM Systems Director.

Syntax

```
smcli sysmirror/help [-l|--list]
smcli sysmirror/help [-v|--verbose] \
  [<COMMAND>[ <COMMAND#2> ...]]
```

Description

The provided help can be general, pertaining to the overall use of the PowerHA SystemMirror commands, or the help can be detailed, specific to one or more commands. All commands have similar operational characteristics.

Most commands accept one or more options that you can use to provide needed or optional information for that command. A short option, indicated with a single leading dash, can be specified with only a single, case-sensitive character, or can be grouped together for convenience (for example, -c or -cSa).

Long options, indicated with a double leading dash, cannot be grouped together, and are not case sensitive. Long options do not always have to be fully typed out. It is only necessary to specify as many leading characters as are required to make the option unique among the other operations available for that operation. Long options are often preferred for scripting purposes, because they help to better document what the script is doing.

Flags

-l|--list

Requests a simple listing of all available commands for PowerHA SystemMirror in IBM Systems Director.

-v|--verbose

Requests maximum details in the displayed information.

Parameter

<COMMAND>

Requests help for the specified command.

importcluster command

Purpose

Imports cluster information into the IBM Systems Director database from a particular node.

Syntax

```
smcli sysmirror/test/importcluster {-h|-?|--help} [-v|--verbose]
smcli sysmirror/test/importcluster [-v|--verbose] <NODE>
```

Command alias

imcl

Description

This command can be used when automatic cluster discovery is not successful. It is not necessary to import cluster data under normal circumstances. It can take a long time to import cluster data.

Flags

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<NODE>

A label that you can use to identify a node.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Example

In the following example, you are importing a cluster that is defined on host `hostA.company.domain.com`:

```
smcli sysmirror/importcluster hostA.company.domain.com
```

Note: You must specify the host name. You cannot use a node name. To identify a host name, use the `smcli sysmirror/lsnode --list` command.

lsappctl command

Purpose

Displays information about the PowerHA SystemMirror application controllers that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lsappctl [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lsappctl {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
[-v|--verbose] \
[<APPCTL>,[<APPCTL#2>,...]]
```

Command alias

lsac

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

-c|--clusters

Displays one or more clusters on which to run this command.

Parameters

<APPCTL>

The label of an application controller to perform this operation on.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Display a simple list of all available application controllers on cluster clusterA:
smcli sysmirror/lsappctl -c clusterA
2. Display the details about appA and appB on cluster clusterA:
smcli sysmirror/lsappctl -c clusterA appA appB
3. Display details about all known application controllers on cluster clusterA:
smcli sysmirror/lsappctl -c clusterA -v

lsappmon command

Purpose

Displays information about the PowerHA SystemMirror application monitors that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lsappmon [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lsappmon {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
                          [-v|--verbose] \
                          [<APPMON>,<APPMON#2>,...]
```

Command alias

lsam

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.
- c|--clusters**
Displays one or more clusters on which to run this command.

Parameters

<APPMON>
The label of an application monitor to perform this operation on.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.

- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Display a simple list of all available monitors on cluster clusterA:
`smcli sysmirror/lsappmon -c clusterA`
2. Display the details about monA and monB on cluster clusterA:
`smcli sysmirror/lsappmon -c clusterA monA monB`
3. Display details about all known application monitors on cluster clusterA:
`smcli sysmirror/lsappmon -c clusterA -v`

Isclrole command

Purpose

Displays information about the security roles that are currently available on the specified PowerHA SystemMirror cluster.

Syntax

```
smcli sysmirror/lsclrole {-h|-?|--help} [-v|--verbose]
smcli sysmirror/lsclrole {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
    [-v|--verbose] \
    [<role_name>[,<role_name#2>,...]]
```

Command alias

lscr

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.
- c|--clusters**
Displays one or more clusters on which to run this command.

Parameters

- <role_name>**
An AIX role on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Display a simple list of all available roles on cluster clusterA:
`smcli sysmirror/lscrole -c clusterA`
2. Display the details about ha_admin on cluster clusterA:
`smcli sysmirror/lscrole -c clusterA ha_admin`
3. Display details about all known roles on cluster clusterA:
`smcli sysmirror/lscrole -c clusterA -v`

Iscluster command

Purpose

Displays information about the PowerHA SystemMirror clusters that are that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lsccluster {-h|-?|--help} \  
                           [-v|--verbose] \  
smcli sysmirror/lsccluster [-v|--verbose] \  
                           [<CLUSTER>[,<CLUSTER#2>,...]]
```

Command alias

lscl

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

-h|-?|--help
Requests help for this command.

-v|--verbose
Requests maximum details in the displayed information.

Parameter

<CLUSTER>
Specifies the label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Display a simple list of all available file collections on cluster clusterA:
smcli sysmirror/lsfilecollection -c clusterA
2. Display the details about Configuration_Files on cluster clusterA:
smcli sysmirror/lsfilecollection -c clusterA Configuration_Files
3. Display details about all known file collections on cluster clusterA:
smcli sysmirror/lsfilecollection -c clusterA -v

Related information:

mkcluster command

Isdependency command

Purpose

Displays information about the PowerHA SystemMirror resource group dependencies that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lsdependency [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lsdependency {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
[-v|--verbose] \
[{-t|--type} <TYPE>] \
[<DEPENDENCY>,<DEPENDENCY#2>,...]
```

Command alias

lsdp

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.
- c|--clusters**
Displays one or more clusters on which to run this command.
- t|--type**
The type of the dependency, from the following set: PARENT_CHILD, START_AFTER, STOP_AFTER, SAME_NODE, SAME_SITE, DIFFERENT_NODES.

Parameters

- <DEPENDENCY>**
The label of a dependency on which to perform this operation.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Display a simple list of all available dependencies on cluster `clusterA`:
`smcli sysmirror/lsdependency -c clusterA`
2. Display the details about `depA` and `depB` on cluster `clusterA`:
`smcli sysmirror/lsdependency -c clusterA depA depB`
3. Display details about all known dependencies on cluster `clusterA`:
`smcli sysmirror/lsdependency -c clusterA -v`

IsefsKeystore command

Purpose

Displays information about the Encrypted File System (EFS) keystore that is configured on the specified PowerHA SystemMirror cluster.

Syntax

```
smcli sysmirror/lsefsKeystore {-h|-?|--help} [-v|--verbose]
smcli sysmirror/lsefsKeystore [-v|--verbose] <CLUSTER>
```

Command alias

lsef

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<CLUSTER>

Specifies the label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Display the EFS keystore, if available, on cluster clusterA:
smcli sysmirror/lsef clusterA

2. Display the details about the EFS configuration on cluster clusterA:

```
smcli sysmirror/lsef -v clusterA
```

lsfallbacktimer command

Purpose

Displays information about fallback timers that are configured on the specified PowerHA SystemMirror cluster.

Syntax

```
smcli sysmirror/lsfallbacktimer {-h|-?|--help} [-v|--verbose]
smcli sysmirror/lsfallbacktimer {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
    [-v|--verbose] \
    [<FALLBACKTIMER>,<FALLBACKTIMER#2>,...]
```

Command alias

lsft

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.
- c|--clusters**
Displays one or more clusters on which to run this command.

Parameters

<FALLBACKTIMER>
The label of a fallback timer on which to perform this operation.

Exit status

- 0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.

- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Display a simple list of all available fallback timers on cluster clusterA:
`smcli sysmirror/lstfallbacktimer -c clusterA`
2. Display the details about myTimer on cluster clusterA:
`smcli sysmirror/lstft -c clusterA myTimer`
3. Display details about all known fallback timers on cluster clusterA:
`smcli sysmirror/lstft -c clusterA -v`

lsfilecollection command

Purpose

Displays information about the PowerHA SystemMirror file collections that are that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lsfilecollection [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lsfilecollection {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
    [-v|--verbose] \
    [<FC>[,<FC#2>,...]]
```

Command alias

lsfc

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.
- c|--clusters**
Displays one or more clusters on which to run this command.

Parameters

- <CLUSTER>**
The label of a cluster on which to perform this operation.
- <FC>**
A file collection on which to perform this operation.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Display a simple list of all available file collections on cluster clusterA:
smcli sysmirror/lfilecollection -c clusterA
2. Display the details about Configuration_Files on cluster clusterA:
smcli sysmirror/lfilecollection -c clusterA Configuration_Files
3. Display details about all known file collections on cluster clusterA:
smcli sysmirror/lfilecollection -c clusterA -v

lsgroup command

Purpose

Displays information about user groups that are configured on the specified PowerHA SystemMirror cluster.

Syntax

```
smcli sysmirror/lsgroup {-h|-?|--help} [-v|--verbose]
smcli sysmirror/lsgroup {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
[-R registry]{ldap|files}
smcli sysmirror/lsgroup {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
[-v|--verbose] \
[<group_name#1>[,<group_name#2>,...]]
```

Command alias

lsgr

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- c|--clusters**
Displays one or more clusters on which to run this command.
- h|-?|--help**
Requests help for this command.

-R|--registry

Indicates where authentication information and authentication credentials are currently stored or where they are going to be stored.

-v|--verbose

Requests maximum details in the displayed information.

Parameters**<group name>**

An AIX user group on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Display a simple list of all available groups on cluster clusterA:
`smcli sysmirror/lsgroup -c clusterA`
2. Display details about all known groups on cluster clusterA:
`smcli sysmirror/lsgr -c clusterA -v`

lsinterface command**Purpose**

Displays information about the PowerHA SystemMirror interfaces that are that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lsinterface [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lsinterface {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
                             [-v|--verbose] \
                             [<INTF>[,<INTF#2>,...]]
```

Command alias

lsif

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.
- c|--clusters**
Displays one or more clusters on which to run this command.

Parameters

- <CLUSTER>**
The label of a cluster on which to perform this operation.
- <INTERFACE>**
The label of an interface on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Display a simple list of all available interfaces on cluster clusterA:
`smcli sysmirror/lsinterface -c clusterA`
2. Display the details about interfaceA and interfaceB on cluster clusterA:
`smcli sysmirror/lsinterface -c clusterA interfaceA interfaceB`
3. Display details about all known interfaces on cluster clusterA:
`smcli sysmirror/lsinterface -c clusterA -v`

lsldapcl command

Purpose

Displays information about the Lightweight Directory Access Protocol (LDAP) client configuration on the specified PowerHA SystemMirror cluster.

Syntax

```
smcli sysmirror/lsldapcl {-h|-?|--help} [-v|--verbose]
smcli sysmirror/lsldapcl [-v|--verbose] <CLUSTER>
```

Command alias

lslc

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<CLUSTER>

Specifies the label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Display the details about an LDAP client configuration, if available, on cluster clusterA:
smcli sysmirror/lsldapcl clusterA

2. Display the details about an LDAP client configuration on cluster clusterA:

```
smcli sysmirror/lslc -v clusterA
```

lsldapsrv command

Purpose

Displays information about the Lightweight Directory Access Protocol (LDAP) server configuration on the specified PowerHA SystemMirror cluster.

Syntax

```
smcli sysmirror/lsldapsrv {-h|-?|--help} [-v|--verbose]
smcli sysmirror/lsldapsrv [-v|--verbose] <CLUSTER>
```

Command alias

lssr

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

-c|--clusters

Displays one or more clusters on which to run this command.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Display the details about an LDAP server configuration, if available, on cluster clusterA:
`smcli sysmirror/lslldapsrv clusterA`
2. Display the details about an LDAP server configuration on cluster clusterA:
`smcli sysmirror/lssr -v clusterA`

Islog command

Purpose

Displays information about the PowerHA SystemMirror logs that exist on the clusters that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lislog [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lislog {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
                        [-v|--verbose] \
                        [<LOG>,<LOG#2>,...]
```

Command alias

lslg

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

-c|--clusters

Displays one or more clusters on which to run this command.

Parameters

<LOG>

The label of a log on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.

- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Display a simple list of all available logs on cluster clusterA:
`smcli sysmirror/l slog -c clusterA`
2. Display the details about logA and logB on cluster clusterA:
`smcli sysmirror/l slog -c clusterA logA logB`
3. Display details about all known logs on cluster clusterA:
`smcli sysmirror/l slog -c clusterA -v`

lsmethod command

Purpose

Displays information about the PowerHA SystemMirror methods that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lsmethod [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lsmethod {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
                        [-v|--verbose] \
                        [{-t|--type} <TYPE>] \
                        [<METHOD>,<METHOD#2>,...]
```

Command alias

lsmd

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.
- c|--clusters**
One or more clusters to run this command on.
- t|--type**
One of the following types of methods: notify, snapshot, verify.

Parameters

- <METHOD>**
The label of a method on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Display a simple list of all available methods, of all types, on cluster clusterA:
`smcli sysmirror/lsmethod -c clusterA`
2. Display the details about methodA and methodB on cluster clusterA:
`smcli sysmirror/lsmethod -c clusterA methodA methodB`
3. Display a particular type of method, such as verify, on clusterA:
`smcli sysmirror/lsmethod -c clusterA -t verify`

lsnode command

Purpose

Displays information about the PowerHA SystemMirror nodes currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lsnode [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lsnode {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
                        [-v|--verbose] \
                        [<NODE>[,<NODE#2>,...]]
smcli sysmirror/lsnode {-a|--available}
```

Command alias

lsnd

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

-h|-?|--help
Requests help for this command.

- v|--verbose**
Requests maximum details in the displayed information.
- c|--clusters**
One or more clusters on which to run this command.
- a|--available**
Indicates all available functions, and displays all unused items of the class.

Parameters

- <CLUSTER>**
The label of a cluster on which to perform this operation.
- <NODE>**
The label of a node on which to perform this operation.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Display a simple list of all available nodes on cluster clusterA:
smcli sysmirror/lsnode -c clusterA
2. Display the details about nodeA and nodeB nodes on cluster clusterA:
smcli sysmirror/lsnode -c clusterA nodeA nodeB
3. Display details about all known nodes on cluster clusterA:
smcli sysmirror/lsnode -c clusterA -v

lsnetwork command

Purpose

Displays information about the PowerHA SystemMirror networks that are that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lsnetwork [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lsnetwork {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
[-v|--verbose] \
[<NETWORK>[,<NETWORK#2>,...]]
```

Command alias

lsnw

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

-c|--clusters

Displays one or more clusters on which to run this command.

Parameters

<CLUSTER>

The label of a cluster on which to perform this operation.

<NETWORK>

The label of a network on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Display a simple list of all available networks on cluster clusterA:
smcli sysmirror/lsnetwork -c clusterA
2. Display the details about networkA and networkB on cluster clusterA:
smcli sysmirror/lsnetwork -c clusterA networkA networkB

lspersistentip command

Purpose

Displays information about the PowerHA SystemMirror persistent IP addresses that are that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lspersistentip {-h|-?|--help} [-v|--verbose]
smcli sysmirror/lspersistentip {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
    [-v|--verbose] \
    [<IP>,[<IP#2>,...]]
```

Command alias

lspi

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.
- c|--clusters**
Displays one or more clusters on which to run this command.

Parameters

- <IP>**
A persistent IP or name on which to perform this operation.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Display a simple list of all persistent IP addresses on cluster clusterA:
smcli sysmirror/lspersistentip -c clusterA
2. Display the details about pipA and pipB on cluster clusterA:
smcli sysmirror/lspersistentip -c clusterA pipA pipB
3. Display details about all known persistent IP addresses on cluster clusterA:
smcli sysmirror/lspersistentip -c clusterA -v

Isresgrp command

Purpose

Lists the available resource groups, or the details about one or more resource groups.

Syntax

```
smcli sysmirror/lresgrp [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lresgrp {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
                        [-v|--verbose] \
                        [<RG>[,<RG#2>,...]]
```

Command alias

lsrg

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- h|-?|--help
Requests help for this command.
- v|--verbose
Requests maximum details in the displayed information.
- c|--clusters
Displays one or more clusters on which to run this command.

Parameters

<RG>
The label of a resource group on which the operation is performed.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.

- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. To display a simple list of all available resource groups on cluster clusterA:
`smcli sysmirror/lresgrp -c clusterA`
2. To display the details about rgA and rgB on cluster clusterA:
`smcli sysmirror/lrsg -c clusterA rgA rgB`
3. To display the details about all known resource groups on cluster clusterA:
`smcli sysmirror/lrsg -c clusterA -v`

Isserviceip command

Purpose

Displays information about the PowerHA SystemMirror service IP addresses that are that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lsserviceip {-h|-?|--help} [-v|--verbose]
smcli sysmirror/lsserviceip {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
    [-v|--verbose] \
    [<IP>,<IP#2>,...]
```

Command alias

lssi

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.
- c|--clusters**
Displays one or more clusters on which to run this command.

Parameters

<IP>
A service IP or name on which to perform this operation.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Display a simple list of all service IP addresses on cluster clusterA:
`smcli sysmirror/lsserviceip -c clusterA`
2. Display the details about sipA and sipB on cluster clusterA:
`smcli sysmirror/lsserviceip -c clusterA sipA sipB`
3. Display details about all known service IP addresses on cluster clusterA:
`smcli sysmirror/lsserviceip -c clusterA -v`

Issite command

Purpose

Displays information about PowerHA SystemMirror sites that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lssite [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lssite {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
                        [-v|--verbose] \
                        [<SITE>,<SITE#2>,...]
```

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.
- c|--clusters**
Displays one or more clusters on which to run this command.

Parameters

<SITE>

Specifies the label of a site on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Display a simple list of all available sites on cluster `clusterA`:
`smcli sysmirror/lssite -c clusterA`
2. Display the details about `siteA` and `siteB` on cluster `clusterA`:
`smcli sysmirror/lssite -c clusterA siteA siteB`
3. Display details about all known sites on cluster `clusterA`:
`smcli sysmirror/lssite -c clusterA -v`

Issnapshot command

Purpose

Displays information about the PowerHA SystemMirror snapshots that are currently managed by the IBM Systems Director.

Syntax

```
smcli sysmirror/lssnapshot [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lssnapshot {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
                             [-v|--verbose] \
                             [<SNAPSHOT>,<SNAPSHOT#2>,...]
```

Command alias

`lsss`

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

-c|--clusters

Displays one or more clusters on which to run this command.

Parameters

<SNAPSHOT>

The label of a snapshot on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Display a simple list of all available snapshots on cluster `clusterA`:
`smcli sysmirror/lssnapshot -c clusterA`
2. Display the details about `snapshotA` and `snapshotB` on cluster `clusterA`:
`smcli sysmirror/lssnapshot -c clusterA snapshotA snapshotB`
3. Display details about all known snapshots on cluster `clusterA`:
`smcli sysmirror/lssnapshot -c clusterA -v`

lstape command

Purpose

Displays information about the PowerHA SystemMirror tapes that are currently managed by the IBM Systems Director server.

Syntax

```
smcli sysmirror/lstape [-h|-?|--help] [-v|--verbose]
smcli sysmirror/lstape {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
    [-v|--verbose] \
    [<TAPE>[,<TAPE#2>,...]]
```

Command alias

lstp

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

-c|--clusters

Displays one or more clusters on which to run this command.

Parameters

<CLUSTER>

The label of a cluster on which to perform this operation.

<TAPE>

The label of a tape on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Display a simple list of all available tapes on cluster clusterA:
smcli sysmirror/lstape -c clusterA
2. Display the details about tapeA and tapeB on cluster clusterA:
smcli sysmirror/lstape -c clusterA tapeA tapeB
3. Display details about all known tapes on cluster clusterA:
smcli sysmirror/lstape -c clusterA -v

lsversion command

Purpose

Displays information about PowerHA SystemMirror plug-ins that are running on either the IBM Systems Director system or on one or more agent systems.

Syntax

```
smcli sysmirror/lsversion [-h|-?|--help]
smcli sysmirror/lsversion [-s|--server] [-v|--verbose]
smcli sysmirror/lsversion [-a|--agents] [-v|--verbose]
smcli sysmirror/lsversion [-a|--agents] [-v|--verbose] \
    {-n|--nodes} <NODE>[,<NODE#2>,...]
smcli sysmirror/lsversion [-a|--agents] [-v|--verbose] \
    {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...]
```

Command alias

lsvr

Description

If verbose information is requested, the build and installation dates of the plug-in is displayed.

Flags

-a|--agents

Indicates that you want to display agent plug-in information. If you do not specify the **lsversion -c** flag or the **lsversion -c** flag, all known agents are queried and displayed.

-c|--clusters

Limits the list of agents that are displayed to those that are running within the specified clusters. Clusters can be specified by their name, label, or ID number.

-h|-?|--help

Requests help for this command.

-n|--nodes

Limits the list of agents that are displayed to those that are running on the specified nodes. Nodes can be specified by their name or label.

-s|--server

Indicates that you want to display server plug-in information. This flag is the default mode.

-v|--verbose

Requests maximum details in the displayed information.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.

- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Display information about all server plug-ins:
`smcli sysmirror/lsversion -v -s`
2. Display the agent plug-in information for all known PowerHA SystemMirror agents:
`smcli sysmirror/lsvr -v -a`
3. Display the agent plug-in information for all agents within clusterA:
`smcli sysmirror/lsvr -av -c clusterA`
4. Display the agent plug-in information for the agents that are running on nodeA and nodeB:
`smcli sysmirror/lsvr -av -n nodeA,nodeB`

lsvolumegroup command

Purpose

Displays information about the volume groups that are currently defined within the specified PowerHA SystemMirror cluster.

Syntax

```
smcli sysmirror/lsvolumegroup {-h|-?|--help} [-v|--verbose]
smcli sysmirror/lsvolumegroup {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
    [-v|--verbose] \
    [<VG>[,<VG#2>,...]]
```

Command alias

lsvg

Description

Lists all configured objects of the specified class, such as nodes and resource groups. If one or more specific object instances is specified, the details for each of those instances are displayed in a <ATTRIBUTE>=<VALUE> format. If verbose output is requested, more details are provided.

Flags

- c|--clusters**
Displays one or more clusters on which to run this command.
- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

- <VG>**
A label that you can use to identify a volume group.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Display a simple list of all volume groups on cluster `clusterA`:
`smcli sysmirror/lsvg -c clusterA`
2. Display the details about a volume group named `myVG` on cluster `clusterA`:
`smcli sysmirror/lsvg -c clusterA myVG`
3. Display details about all known volume groups on cluster `clusterA`:
`smcli sysmirror/lsvg -c clusterA -v`

mkappctl command

Purpose

Creates a PowerHA SystemMirror application controller definition that conforms to the provided specifications.

Syntax

```
smcli sysmirror/mkappctl [-h|-?|--help] [-v|--verbose]
smcli sysmirror/mkappctl {-c|--cluster} <CLUSTER> \
    {-s|--start_script} <PATH_TO_SCRIPT> \
    {-S|--stop_script} <PATH_TO_SCRIPT> \
    [{-m|--monitors} <MONITOR>,[<MONITOR#2>,...]] \
    <appctl_name>
```

Command alias

`mkac`

Description

This new definition initially exists only on the controlling node of the specified cluster; none of the other nodes in the cluster are aware of it yet. To complete the creation, you need to synchronize the cluster. This is done as a separate step so that you can make multiple changes to the cluster, and then verify those changes as a set before enabling them by synchronization. You can discard your changes at any time before the synchronization occurs, as long as you have not changed your cluster's controlling node.

Flags

- c|--cluster**
The cluster on which to run this command.
- h|-?|--help**
Requests help for this command.
- m|--monitors**
One or more monitors that are to be assigned to this application.
- s|--start_script**
The executable file that is used to start the application.
- S|--stop_script**
The executable file that is used to stop the application.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

- <appctl_name>**
A label by which the application controller can be referred to.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Create an application controller named app1 on clusterA:

```
smcli sysmirror/mkappctl -c clusterA \  
                        -s /usr/local/bin/start_app.sh \  
                        -S /usr/local/bin/stop_app.sh \  
                        app1
```
2. Create an application controller named app2 on clusterA, specifying two monitors:

```
smcli sysmirror/mkappctl -c clusterA \  
                        -s /usr/local/bin/start_app.sh \  
                        -S /usr/local/bin/stop_app.sh \  
                        -m appmon#1,appmon#2 \  
                        app2
```


mkappmon command

Purpose

Creates a PowerHA SystemMirror application monitor definition that conforms to the provided specifications.

Syntax

```
smcli sysmirror/mkappmon [-h|--help] [-v|--verbose]
smcli sysmirror/mkappmon {-c|--cluster} <CLUSTER> \
  {-t|--type} process \
  {-a|--applications} <APPCTL>[,<APPCTL#2>,...] \
  {-m|--mode} {continuous|startup|both} \
  {-p|--processes} <PROC>[,<PROC#2>,...] \
  {-o|--owner} <PROCESSES_OWNER_ID> \
  [{-i|--instance_count} ###] \
  [{-s|--stabilization_interval} ###] \
  [{-r|--restart_count} ###] \
  [{-R|--restart_interval} ###] \
  [{-f|--failure_action} {notify|failover}] \
  [{-N|--notify_method} <PATH_TO_FILE>] \
  [{-C|--cleanup_method} <PATH_TO_FILE>] \
  [{-S|--restart_method} <PATH_TO_FILE>] \
  <appmon_name>
smcli sysmirror/mkappmon {-c|--cluster} <CLUSTER> \
  {-t|--type} custom \
  {-a|--applications} <APPCTL>[,<APPCTL#2>,...] \
  {-m|--mode} {continuous|startup|both} \
  {-M|--monitor_method} <PATH_TO_FILE>] \
  [{-I|--monitor_interval} ###] \
  [{-k|--hung_signal} #] \
  [{-s|--stabilization_interval} ###] \
  [{-r|--restart_count} ###] \
  [{-R|--restart_interval} ###] \
  [{-f|--failure_action} {notify|failover}] \
  [{-N|--notify_method} <PATH_TO_FILE>] \
  [{-C|--cleanup_method} <PATH_TO_FILE>] \
  [{-S|--restart_method} <PATH_TO_FILE>] \
  <appmon_name>
```

Command alias

mkam

Description

Creates an application monitor definition within the specified cluster.

It is important to understand that this new definition will initially exist only on the controlling node of the specified cluster; none of the other nodes in the cluster will be aware of it yet. In order to complete the creation, you will need to synchronize the cluster. This is done as a separate step to allow you to make multiple changes to the cluster, then verify those changes as a set before enabling them via synchronization. You may discard your changes at any time before the synchronization occurs, as long as you have not changes your cluster's controlling node.

Flags

-a|--applications

The label of one or more application controllers to attach this monitor to.

-C|--cleanup_method

The absolute path to a file that is to be run when a failed application is detected, before calling the

restart method. The default is the application controller's stop script. If you are configuring a monitor that is used only in the startup monitoring mode, the method specified in this field does not apply, and will be ignored.

-c|--cluster

The cluster on which to run this command.

-f|--failure_action

The action to be taken if the application cannot be restarted within the specified restart count. Choices are **notify** (the default value) and **fallover**. If **notify** is selected, no further action is taken after running the notify method. If **fallover** is selected, the resource group containing the monitored application is moved to another node in the cluster.

If you are configuring a monitor that is to be used only in the startup monitoring mode, the method specified in this field does not apply, and will be ignored.

-h|--help

Requests help for this command.

-k|--hung_signal

The signal sent to stop the monitor method if it does not complete its processing within the number of seconds established by the monitor interval setting. The default is 9 (SIGKILL).

-i|--instance_count

The number of instances of a process to monitor. The default is 1. The number of instances must exactly match the number of processes to be monitored. If you start one instance, and another instance of the application starts, you will receive an application monitor error. This number must be more than 1 if you have specified more than one process to monitor (1 instance for each process).

-I|--monitor_interval

The monitor method to be run periodically at this interval, specified in seconds. If a monitor method takes longer than this interval to determine the status of its application, the monitor method is stopped by using the specified **hung_signal** flag.

-m|--mode

The manner in which this monitor is run, specified from the following set:

STARTUP

Verifies that the application controller has successfully started within the specified stabilization interval. Select this mode if you are configuring an application monitor for an application that is included in a parent resource group (in addition to other monitors that you might need for dependent resource groups).

CONTINUOUS

Periodically checks that the application controller is running. The checking starts after the specified stabilization interval has passed. This mode is the default.

BOTH Checks that the application controller successfully started within the specified stabilization interval, and then periodically checks that the application controller is still running after the stabilization interval has passed.

-M|--monitor_method

The absolute path to a file that is to be run for custom monitoring of the health of the specified application (started by the application controller).

The method must return a zero to indicate that the application it is monitoring is healthy, or a nonzero value if a problem is detected.

You can have the monitor log messages to the log file `/var/hacmp/log/clappmond.application_controller_name.resource_group_name.lo` by having it print messages to STDOUT. The messages are automatically redirected to the monitor log.

-N|--notify_method

The absolute path to a file that is to be run when a failed application is detected to perform notification. This method runs each time an application is restarted, fails completely, or falls over to the next node in the cluster. If you are configuring a monitor that is going to be used only in the startup monitoring mode, the method specified in this field does not apply, and will be ignored.

-p|--processes

Specify one or more processes to monitor. Ensure that you are using correct process names, by using the names as they appear in the output of the **ps -el** command.

-o|--owner

Specify the user ID of the owner of the processes that are being monitored, for example, root. This owner must own all processes that are being monitored.

-S|--restart_method

The absolute path to a file that is to be processed when a failed application is detected to attempt to restart the application. The default is the application controller's start script. If you are configuring a monitor that is to be used only in the startup monitoring mode, the method specified in this field does not apply, and will be ignored.

-r|--restart_count

The number of times to try restarting the application before taking any other actions. The default value is 3. If you are configuring a monitor that is going to be used only in the startup monitoring mode, the restart count does not apply, and will be ignored. Make sure that you also provide a restart method if a restart count is specified.

-R|--restart_interval

The number of seconds that the application must remain stable before resetting the restart count. Do not set this value to be shorter than **Restart_Count multiplied by Stabilization_Interval**. The default is 10% longer than that value. If the restart interval is too short, the restart count is reset too soon and the desired fallover or notify action might not occur when it should. If you are configuring a monitor that is to be used only in the startup monitoring mode, the restart interval does not apply, and will be ignored.

-s|--stabilization_interval

The number of seconds to wait for the application to stabilize. This value is used in different ways, depending on which mode is specified for this monitor. In most circumstances, this value should not be zero.

STARTUP

The stabilization interval is the period within which the monitor is run to check that the application successfully started. When the specified time ends, PowerHA SystemMirror stops the monitoring of the application startup and continues event processing.

If the application fails to start within the stabilization interval, the resource group's acquisition fails on the node, and PowerHA SystemMirror launches resource group recovery actions to acquire the resource group on another node. The number of seconds you specify should be approximately equal to the period of time it takes for the application to start. This depends on the application you are using.

CONTINUOUS

The stabilization interval is the period during which PowerHA SystemMirror waits before beginning to check that the application is running successfully. For instance, with a database application, you might want to delay monitoring until after the start script and initial database search have been completed. You might need to experiment with this value to balance performance with reliability.

BOTH The stabilization interval is used to wait for the application to start successfully, and then is used to wait until it starts checking periodically that the application is successfully running on the node.

-t|--type
The type of the application monitor: Process or Custom.

-v|--verbose
Requests maximum details in the displayed information.

Parameters

<appmon_name>
A label by which the application monitor can be referred to.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Create an application process monitor named procmon that performs continuous monitoring:

```
smcli sysmirror/mkappmon -c clusterA \  
    -t process \  
    -m continuous \  
    -a db2app \  
    -p db2procname \  
    -o db2procowner \  
    procmon
```

2. Create a custom application monitor named usermon that monitors the startup of an application:

```
smcli sysmirror/mkappmon -c clusterA \  
    -t custom \  
    -m startup \  
    -a appctlr \  
    -M /file/to/run/for/monitoring \  
    usermon
```

mkcluster command

Purpose

Creates a PowerHA SystemMirror cluster definition that conforms to the provided specifications.

Syntax

```
smcli sysmirror/mkcluster {-h|-?|--help} [-v|--verbose]  
smcli sysmirror/mkcluster {-c|--controlling_node} <node> \  
    [{-i|--cluster_ip} <multicast_address>] \  
    [{-t|--type} {NSC|SC|LC} ] \  
    \  
    <nsd> <scd> <lcsd> <nsdscd> <nsdscld> <nsdscldscd> <nsdscldscldscd>
```

```
[{-r|--repositories} <disk>[,<backup#1,...>] ] \
[{-n|--nodes} <NODE>[,<NODE#2>,...] ] \
[{-S|--fc_sync_interval} <##>] \
[{-s|--rg_settling_time} <##>] \
[{-e|--max_event_time} <##>] \
[{-R|--max_rg_processing_time} <##>] \
[{-D|--disable_daily_verification} ] \
[{-V|--verification_node} {Default|<NODE>} ] \
[{-H|--verification_hour} <00..23>] \
[{-m|--minimize_ver_output} ] \
[{-F|--site_policy_failure_action} {fallover|notify}] \
[{-N|--site_policy_notify_method} <path_to_file>] \
[{-g|--site_grace_period} <period>] \
[<cluster_name>]]
```

Command alias

mkcl

Description

Creates a cluster definition on the specified controlling node. If no controlling node is specified, one of the nodes from the node list is arbitrarily chosen.

This new definition initially exists only on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of it yet. To complete the creation process, the cluster must be synchronized by using the **synccluster** command.

This synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

-h|-?|--help

Requests help for this command.

-i|--comppath

The communication path to take to the node, such as an IP address or host name.

-c|--controlling_node

The node that serves as the primary point of contact in this cluster for the Director server.

-D|--disable_daily_verification

Disables the automatic daily verification of the cluster configuration. You do not have to use this flag under normal circumstances.

-e|--max_event_time

The maximum amount of time that an event can run before reporting a config_too_long error.

-F|--site_policy_failure_action

The action to take when a resource group cannot be brought online from its primary site.

-g|--site_grace_period

The site grace period in seconds. This flag is only available for linked clusters (LC) and stretched clusters (SC) that are within sites.

-H|--verification_hour

The hour when the automatic, daily verification of the cluster configuration is invoked. The range is 00 - 23, where 00 indicates that the time is midnight and 23 indicates that the time is 11:00 p.m.

-i|--cluster_ip

The multicast IP address to be used by the cluster.

- m|--minimize_ver_output**
Indicates that no debug output is required during the automatic daily verification of the cluster configuration.
- n|--nodes**
A list of nodes that are the members of the cluster.
- N|--site_policy_notify_method**
The notification script to run when the site policy action is set to notify.
- r|--repository**
The shared, central repository database disk for the cluster.
- R|--max_rg_processing_time**
The maximum amount of time to allow a resource group to be acquired or released before reporting a config_too_long error.
- s|--rg_settling_time**
The number of seconds to wait for resource groups to initialize before actively checking on their status.
- S|--fc_sync_interval**
The number of seconds between cluster polling to determine whether any file collections have to be synchronized.
- t|--type**
The type of cluster. You can specify one of the following types of a cluster:
 - NSC (no site)
 - SC (stretched)
 - LC (linked)
- v|--verbose**
Requests maximum details in the displayed information.
- V|--verification_node**
The node that is used for the automatic daily verification of the cluster configuration. If you specify a value of *Default*, the cluster selects the node automatically.

Parameters

<cluster_name>
A label by which the cluster can be referred to.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.

29 The current locale is not supported.

Examples

1. Define a standard no-sites cluster named `clusterA` that incorporates nodes `nodeA` and `nodeB`. The cluster uses repository disk `hdisk3`, and a multicast address of `227.1.2.3`:

```
smcli sysmirror/mkcluster -t NSC \  
    -n nodeA,nodeB \  
    -r hdisk3 \  
    -i 227.1.2.3 \  
    clusterA
```

2. Define a stretched cluster named `clusterB` that incorporates nodes `nodeA`, `nodeB`, `nodeC`, and `nodeD`. The cluster uses repository disk `hdisk6`, and a multicast address of `228.4.5.6`:

```
smcli sysmirror/mkcluster -t SC \  
    -n nodeA,nodeB,nodeC,nodeD \  
    -r hdisk6 \  
    -i 228.4.5.6 \  
    clusterB
```

Note: After you run this command, you must run the `mksite` command to define the sites in the stretched cluster.

3. Define a linked cluster named `clusterC` that incorporates nodes `nodeA`, `nodeB`, `nodeC`, and `nodeD`:

```
smcli sysmirror/mkcluster -t SC \  
    -n nodeA,nodeB,nodeC,nodeD \  
    clusterC
```

Note: After you run this command, you must run the `mksite` command to define the sites in the linked cluster and in the repository disks, for each site.

mkdependency command

Purpose

Creates a PowerHA SystemMirror resource group that has a dependency definition that conforms to specifications that you identify.

Syntax

```
smcli sysmirror/mkdependency [-h|-?|--help] [-v|--verbose]
```

Parent/Child Dependency:

```
smcli sysmirror/mkdependency {-c|--cluster} <CLUSTER> \  
    {-p|--parent} <RG> \  
    {-C|--children} <RG>[,<RG#2>,...]
```

Start After Dependency:

```
smcli sysmirror/mkdependency {-c|--cluster} <CLUSTER> \  
    {-s|--start} <RG>[,<RG#2>,...] \  
    {-a|--after} <RG>
```

Stop After Dependency:

```
smcli sysmirror/mkdependency {-c|--cluster} <CLUSTER> \  
    {-S|--stop} <RG>[,<RG#2>,...] \  
    {-a|--after} <RG>
```

Same Node Dependency:

```
smcli sysmirror/mkdependency {-c|--cluster} <CLUSTER> \  
    {-n|--same_node} \  
    {-g|--groups} <RG>,<RG#2>,...
```

Same Site Dependency:

```
smcli sysmirror/mkdependency {-c|--cluster} <CLUSTER> \  
    {-t|--same_site} \  
    {-g|--groups} <RG>,<RG#2>,...
```

Different Nodes Dependency:

```
smcli sysmirror/mkdependency {-c|--cluster} <CLUSTER> \  
                              {-H|--high} <RG>,<RG#2>,... \  
                              {-i|--intermediate} <RG>,<RG#2>,... \  
                              {-l|--low} <RG>,<RG#2>,...
```

Resource Group Processing Order:

```
smcli sysmirror/mkdependency {-c|--cluster} <CLUSTER> \  
                              {-A|--acquire} \  
                              {-r|--serial} <RG>,<RG#2>,...  
smcli sysmirror/mkdependency {-c|--cluster} <CLUSTER> \  
                              {-A|--acquire} \  
                              {-P|--parallel} <RG>,<RG#2>,...
```

```
smcli sysmirror/mkdependency {-c|--cluster} <CLUSTER> \  
                              {-R|--release} \  
                              {-r|--serial} <RG>,<RG#2>,...
```

```
smcli sysmirror/mkdependency {-c|--cluster} <CLUSTER> \  
                              {-R|--release} \  
                              {-P|--parallel} <RG>,<RG#2>,...
```

Command alias

mkdp

Description

Use this command to create a dependency between two or more resource groups, which indicates where the dependency is run (also known as location dependencies) or when the dependency ran (also known as temporal dependencies). Location dependencies focus on restricting resource groups to which nodes a dependency can run on. Temporal dependencies control the order in which resource groups start and stop.

A new dependency initially exists only on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of the new dependency. To complete the creation of the dependency, you must synchronize the cluster. The synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes as a set before enabling them by using synchronization. You can discard your changes at any time before the synchronization occurs if you have not changed the controlling node in the cluster.

Flags

-A|--acquire

Indicates the order in which resource groups are acquired.

-a|--after

Indicates a resource group that you want to be started or stopped before any associated target resource groups.

-C|--children

Indicates one or more resource groups that are to be started after, and stopped before, any associated parent resource group.

-c|--clusters

Displays one or more clusters on which to run this command.

-g|--groups

Indicates one or more resource groups that run on the same node or site.

-h|-?|--help

Requests help for this command.

-H|--high

Sets up a dependency where the resource groups are going to run on different nodes. High priority resource groups are started before lower priority resource groups, and they cannot be on the same node. During a fallback or a fallover, high priority resource groups take precedence and are processed before lower priority resource groups. If a high priority resource group is acquired on a node where a lower priority resource group is already running, the lower priority resource group is released in favor of bringing the high priority resource group online on that node.

-i|--intermediate

Sets up a dependency where the resource groups are going to run on different nodes. Intermediate priority resource groups are started on the first available node that does not currently have a high priority resource group running. During a fallback or a fallover, resource groups with intermediate priority are handled after the high priority resource groups, but before the low priority resource groups. If an intermediate priority resource group is acquired on a node where a low priority resource group is already running, the low priority resource group is released in favor of bringing the intermediate priority resource group online on that node.

-l|--low

Sets up a dependency where the resource groups are going to run on different nodes. Resource groups with low priority are acquired on nodes where no higher priority resource groups are already running. During a fallback or a fallover, the resource groups with high and intermediate priority takes precedence. If a higher priority resource group is acquired on a node, the low priority resource group is moved, or taken offline, and the higher priority resource group is brought online in its place. When acquired, resource groups with equal priority cannot come online on the same node.

-n|--same_node

Indicates that this dependency results in all associated resource groups that are running on the same node.

-P|--parallel

Specifies the set of resource groups that are to be released or acquired at the same time.

-p|--parent

Specifies a resource group that starts before, and stops after, its associated child resource groups.

-R|--release

Modifies the order in which resource groups are released.

-r|--serial

Specifies the set of resource groups that are released or acquired one at a time, one after the other, and in the specified left to right order.

-s|--start

Specifies one or more resource groups that are started after their associated source resource group.

-S|--stop

Specifies one or more resource groups that are stopped after their associated source resource group.

-t|--same_site

Indicates that all associated resource groups are running on the same site.

-v|--verbose

Requests maximum details in the displayed information.

Parameters**<RG>**

A resource group that is involved in the dependency relationship.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Create a parent and child dependency between RGparent, RGchild1, and RGchild2 on cluster clusterA:
smcli sysmirror/mkdependency -c clusterA -p RGparent -C RGchild1, RGchild2
2. Create a dependency on cluster clusterA that verifies that RG1 and RG2 are always running on the same node:
smcli sysmirror/mkdp -c clusterA --same_node --groups RG1, RG2
3. Create a dependency on cluster clusterA that verifies that RG1 and RG2 are never running on the same node:
smcli sysmirror/mkdp -c clusterA --high RG1 --low RG2

mkefsKeystore command

Purpose

Defines an encrypted file system (EFS) keystore on the specified PowerHA SystemMirror cluster.

Syntax

```
smcli sysmirror/mkefsKeystore {-h|-?|--help} [-v|--verbose]
smcli sysmirror/mkefsKeystore {-m|--mode} ldap \
[-p|--efs_password] \
<CLUSTER>
smcli sysmirror/mkefsKeystore {-m|--mode} local \
[-p|--efs_password] \
{-V|--keystore_vg} <VOLUME_GROUP> \
{-s|--service_ip} <SERVICE_IP> \
<CLUSTER>
```

Command alias

mkef

Description

The EFS keystore definition initially exists only on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of the EFS keystore definition. To complete the creation of the EFS keystore definition, you must synchronize the cluster. The synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes. You can discard your changes at any time before the synchronization occurs if you have not changed the controlling node in the cluster.

Flags

-h|-?|--help

Requests help for this command.

-m|--mode

Select the location of the keystore. For example, the Lightweight Directory Access Protocol (LDAP) server or a shared file system.

-p|--password

The keystore admin password, which is used to manage EFS by allowing keys and permissions to be managed.

-s|--service_ip

The name of a service IP address that is used for mounting a shared network file system (NFS).

-V|--keystore_vg

The name of a volume group that will store the keystore file system.

-v|--verbose

Requests maximum details in the displayed information.

Parameter

-c|--clusters

Displays one or more clusters on which to run this command.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Create an EFS keystore in an LDAP server that is defined for cluster clusterA:
smcli sysmirror/mkefsKeystore -m ldap -p myPassword clusterA
2. Create an EFS keystore in the shared file system that is defined within the EFS_VG volume:
smcli sysmirror/mkefsKeystore -m local -p myPassword -v EFS_VG clusterA

mkfallbacktimer command

Purpose

Creates a PowerHA SystemMirror delayed fallback timer definition.

Syntax

```
smcli sysmirror/mkfallbacktimer [-h|-?|--help] [-v|--verbose]
smcli sysmirror/mkfallbacktimer {-c|--cluster} <CLUSTER> \
    {-r|--repeats} <{0..4}|daily|weekly|
    monthly|yearly|once> \
    [{-y|--year} <####>] \
    [{-w|--week_day} <{0..6}|Sun..Sat>] \
    [{-m|--month} <{1..12}|Jan..Dec>] \
    [{-d|--day} {1..31}] \
    {-H|--hour} {0..23} \
    {-M|--minute} {0..59} \
    <FALLBACK_TIMER>
```

Command alias

mkft

Description

Creates a fallback timer definition within the specified cluster that you can assign to a resource group to delay when it falls back to a prior managing node. This process is only available if the current fallback policy for the specified resource group is set to Fallback To Higher Priority Node.

The fallback timer definition initially exists only on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of the fallback timer definition. To create a fallback timer definition, you must synchronize the cluster. This synchronization is completed as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

- c|--clusters**
Displays one or more clusters on which to run this command.
- d|--day**
Specify the day of the fallback timer execution time.
- H|--hour**
Specify the hour of the fallback timer execution time.
- h|-?|--help**
Requests help for this command.
- r|--repeats**
Specify the type and frequency of the fallback timer.
- m|--month**
Specify the month of the fallback timer execution time.
- M|--minute**
Specify the minute of the fallback timer execution time.
- v|--verbose**
Requests maximum details in the displayed information.
- w|--week_day**
Specify the day of the week that corresponds to the fallback timer execution time.
- y|--year**
Specify the year of the fallback timer execution time.

Parameters

<FALLBACK_TIMER>

A label that you can use to reference the fallback timer.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Create a daily fallback timer named myTimer on cluster clusterA, which tries to fallback to the prior managing node (the higher priority node) once per day at 3:30 a.m.:

```
smcli sysmirror/mkfallbacktimer -c clusterA -r daily -H 3 -M 30
```
2. Create a monthly fallback timer named myTimer on cluster clusterA, which tries to fallback to the prior managing node (the higher priority node) once per month during a maintenance time period at 3:30 a.m. on the first day of the month:

```
smcli sysmirror/mkft -c clusterA -r monthly -H 3 -M 30 -d 1
```

mkfilecollection command

Purpose

Creates a PowerHA SystemMirror file collection that conforms to the provided specifications.

Syntax

```
smcli sysmirror/mkfilecollection [-h|-?|--help] [-v|--verbose]
smcli sysmirror/mkfilecollection {-c|--cluster} <CLUSTER> \
    [{-C|--sync_with_cluster}] \
    [{-M|--sync_when_changed}] \
    [{-d|--description} "<DESCRIPTION>"] \
    [{-f|--files} <FILE>[,<FILE#2>,...] ] \
    <file_collection_name>
```

Command alias

mkfc

Description

Creates a file collection within the specified cluster. A file collection can be used to keep one or more files and directories synchronized on all nodes in the cluster.

It is important to understand that this new definition will initially exist only on the controlling node of the specified cluster; none of the other nodes in the cluster will be aware of it yet. In order to complete the creation, you will need to synchronize the cluster. This is done as a separate step to allow you to make multiple changes to the cluster, then verify those changes as a set before enabling them using synchronization. You can discard your changes at any time before the synchronization occurs, as long as you have not changes your cluster's controlling node.

Flags

-c|--cluster

The cluster on which to run this command.

-C|--sync_with_cluster

Indicates that this file collection should automatically be synchronized every time the cluster itself is synchronized.

-d|--description

A brief, one-line sentence that describes the object.

-f|--files

One or more files or directories to include in the file collection. If a director is specified, all of its files and child directories are included in a recursive manner.

-h|-?|--help

Requests help for this command.

-M|--sync_when_changed

Indicates that this file collection should automatically be synchronized every time the cluster detects that one of the collection's files has been changed. This detection process occurs every 10 minutes, but can be adjusted.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<file_collection_name>

A label by which the file collection is referred to.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Create an empty file collection named myfc:

```
smcli sysmirror/mkfilecollection -c clusterA \  
myfc
```

2. Create a file collection that contains two files:

```
smcli sysmirror/mkfilecollection -c clusterA \  
--files "/my/file,/your/file" \  
myfc
```

3. Create a file collection with maximum, automatic synchronization:

```
smcli sysmirror/mkfilecollection -c clusterA \  
--files "/etc/passwd,/etc/group" \  
--sync_with_cluster \  
--sync_when_changed \  
myfc
```

mkgroup command

Purpose

Creates an AIX user group that conforms to the provided specifications.

Syntax

```
smcli sysmirror/mkgroup {-h|-?|--help} [-v|--verbose]  
smcli sysmirror/mkgroup {-c|--cluster} <CLUSTER> \  
{-R|--registry} {ldap|files}} \  
[{-r|--resgrp} <resource_group>] \  
[{-i|--id} #####] \  
[-A|--administrative] \  
[{-u|--users} <USER>[,<USER#2>,...]] \  
[{-a|--admins} <ADMIN>[,<ADMIN#2>,...]] \  
[{-p|--projects} <PROJECT>[,<PROJECT#2>,...]] \  
[{-m|--keystore_mode} {admin|guard}] \  
[{-k|--keystore_encryption} {RSA_1024|RSA_2048|  
RSA_4096}] \  
[{-k|--keystore_access} {ldap|file}] \  
<group_name>
```

Command alias

mkgr

Description

You can use the **mkgroup** command to create an AIX user group on all the nodes in the cluster or on all the nodes in a resource group.

Flags

-A|--administrative

Indicates whether the group is an administrative group. Only the root user can modify the attributes of an administrative group.

-a|--admins

Specify the group members that can work with the attributes for the group. For example, you can add user to the group or remove users from the group. The group attributes of an administrative group can be modified by only the root user. If the group is an administrative group that you specified with the **-a|--administrative** flag, you cannot define other administrators.

-c|--clusters

Displays one or more clusters on which to run this command.

-h|--help

Requests help for this command.

-i|--id

Specify a unique decimal integer string to associate with the group account on the system. If you do not specify a group ID, it is specified automatically. The specified group ID is created on all nodes in the cluster or it is created in the LDAP server.

Note: If you do not specify this value, the AIX operating system can assign different group IDs for each node.

-K|--keystore_encryption

Specify the algorithm that is used to generate the key for the private key of the group during the keystore creation. The default encryption used is RSA_1024.

-k|--keystore_access

Specify the database type of the keystore.

-m|--keystore_mode

Select one of the following modes for resetting the keystore password for users:

admin A root user, or other security privileged system users, can open the keystore by using the admin key. This user can use the admin key to reset the specified user's password. This user does not have access to the specified user's encrypted files. This flag is the default mode.

guard A root user, or other security privileged system users, cannot open the keystore. These users cannot reset a specified user's keystore password.

-p|--projects

Defines the list of projects that you can assign to a group process. The value is a list of comma-separated project names and is evaluated from left to right. The project name must be a valid project name that is defined in the system. If an invalid project name is found on the list, an error is reported.

-R|--registry

Indicates where authentication information and credentials are, or will be, stored.

-r|--resgrp

Specify a resource group that creates a group on only the nodes that are defined within that resource group.

-u|--users

Specify the names of the users that belong to this group. The members of a group can read, write, and execute a resource or file that is owned by another member of the group that is specified by the access control list of a resource.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<group_name>

A label that you can use to reference the group.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

1 The command has not been used in the intended manner.

- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Create a group named myGroup on cluster clusterA that is stored in the local file system:

```
smcli sysmirror/mkgroup -c clusterA -R files myGroup
```
2. Create a group named yourGroup on cluster clusterA that is stored in LDAP, adding users bart and lisa to the group:

```
smcli sysmirror/mkgroup -c clusterA -R ldap -u bart,lisa yourGroup
```

mkinterface command

Purpose

Creates a PowerHA SystemMirror interface definition that conforms to the provided specifications.

Syntax

```
smcli sysmirror/mkinterface [-h|-?|--help] [-v|--verbose]
smcli sysmirror/mkinterface {-c|--cluster} <CLUSTER> \
                             {-n|--network} <NETWORK> \
                             [{-N|--node} <NODE>] \
                             [{-t|--type} {ether|infiniband}] \
                             [{-i|--interface} <interface>] \
                             <interface_name>
```

Command alias

mkif

Description

Creates an interface definition within the specified cluster.

The new interface definition exists only on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of it yet. To complete the creation, you must synchronize the cluster. This synchronization is completed as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

- c|--cluster**
The cluster on which to run this command.
- h|-?|--help**
Requests help for this command.

- i|--interface**
The interface to be used for the interface definition. For example en4, en5, and en6.
- n|--network**
The PowerHA SystemMirror network that manages the interface.
- N|--node**
The node to which this interface belongs.
- t|--type**
The type of the interface. Valid values are ether or infiniband.

Parameters

- <interface_name>**
A label by which the interface can be referred to.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Create an interface named myint without specifying an interface type:

```
smcli sysmirror/mkinterface -c clusterA \  
                             -n net_ether_01 \  
                             myint
```
2. Create an interface named yourint with a specifying interface type:

```
smcli sysmirror/mkinterface -c clusterA \  
                             -n network2 \  
                             -t ether \  
                             yourint
```

mkldapcl command

Purpose

Configures an LDAP client on all the nodes in the cluster.

Syntax

```
smcli sysmirror/mkldapcl {-h|-?|--help} [-v|--verbose]  
smcli sysmirror/mkldapcl {-c|--cluster} <CLUSTER> \  
                          {-b|--bind_dn} <bind_dn> \  
                          {-P|--password} <LDAP_bind_password> \  
                          ...
```

```
{-B|--base_dn} <Suffix/Base_dn> \
{-k|--key} <full_path_to_SSL_key> \
{-K|--ssl_password} <SSL_key_password> \
{-t|--port} {636|###} \
<server>[<,server#2>..]
```

Command alias

mklc

Description

Configures an LDAP client on all the nodes in the cluster that is based on the LDAP servers that have already been configured.

Flags

- B|--base_dn**
Specify the suffix or the base distinguished name (DN) to search on the LDAP server for users, groups, and other network information entities. The default value is `cn=aixdata,o=ibm`.
- b|--bind_dn**
Specify the LDAP bind DN used to bind to the LDAP server. The DN that you specify must exist on the LDAP server. The ability to perform operations on entries in the LDAP server database from the LDAP client is dependent on the access permissions that are granted to the bind DN on the LDAP server. For example, `cn=admin cn=proxy,o=ibm cn=user,ou=people,cn=aixdata`. The default value is `cn=admin`.
- c|--clusters**
Displays one or more clusters on which to run this command.
- h|-?|--help**
Requests help for this command.
- K|--ssl_password**
Specify the password for the Secure Sockets Layer (SSL) key.
- k|--key**
Specify the full path to the SSL key database, which must be in `.kdb` format. If you do not specify the path, an SSL key database is created automatically.
- P|--password**
Specify the text-only LDAP bind password for the DN used to bind to the LDAP server. The password that you use must match the password on the LDAP server for the specified DN.
- t|--port**
Specify the server port number that is used for communicating with the LDAP server. The default port number is 636.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

<server>

The resolvable host name or IP address of an LDAP server that has already been configured for use in the cluster.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Example

The following is an example of a configuration that uses the LDAP clients on clusterA. clusterA uses an LDAP server called myLDAP.

```
smcli sysmirror/mkldapcl -c clusterA -b cn=admin -P myPassword \  
-B "cn=aixdata,o=ibm" \  
-k /keys/my_ssl_key -K my_ssl_password \  
myLDAP
```

mkldapsrv command

Purpose

Configures one or more LDAP servers for use by the specified cluster.

Syntax

```
smcli sysmirror/mkldapsrv {-h|-?|--help} [-v|--verbose]
```

The following syntax is used for defining a new LDAP configuration:

```
smcli sysmirror/mkldapsrv {-c|--cluster} <CLUSTER> \  
{-a|--admin_dn} <LDAP_admin_dn> \  
{-p|--password} <LDAP_admin_password> \  
{-B|--base_dn} <Suffix/Base_dn> \  
{-k|--key} <full_path_to_SSL_key> \  
{-K|--ssl_password} <SSL_key_password> \  
[{-t|--port} <636|###>] \  
[{-S|--schema} <schema_type>] \  
[{-V|--version} <version>] \  
[{-d|--db2_password} <PASSWORD>] \  
[{-e|--encryption_seed} <SEED>] \  
<server>[,<server#2>,...]
```

The following syntax is used for using an existing LDAP configuration:

```
smcli sysmirror/mkldapsrv {-c|--cluster} <CLUSTER> \  
{-b|--bind_dn} <bind_dn> \  
{-P|--password} <LDAP_bind_password> \  
{-B|--base_dn} <Suffix/Base_dn> \  
{-k|--key} <full_path_to_SSL_key> \  
{-K|--ssl_password} <SSL_key_password> \  
[{-t|--port} <636|###>] \  
<server>[,<server#2>,...]
```

Command alias

mksr

Description

You can use the **mkldapsrv** command to configure a cluster with a pre-configured LDAP server, or to create a configuration on an LDAP host server.

Flags

- a|--admin_dn**
Specify the LDAP administrator distinguished name (DN) used to bind to the LDAP server. The DN you specify must exist on the LDAP server. The ability to perform operations on entries in the LDAP server database from the LDAP client is dependent on the access permissions that are granted to the bind DN on the LDAP server. For example, `cn=admin cn=proxy,o=ibm cn=user,ou=people,cn=aixdata`. Default value for this flag is `cn=admin`.
- B|--base_dn**
Specify the suffix or the base distinguished name (DN) to search on the LDAP server for users, groups, and other network information entities. The default value is `cn=aixdata,o=ibm`.
- b|--bind_dn**
Specify the LDAP bind DN used to bind to the LDAP server. The DN that you specify must exist on the LDAP server. The ability to perform operations on entries in the LDAP server database from the LDAP client is dependent on the access permissions that are granted to the bind DN on the LDAP server. For example, `cn=admin cn=proxy,o=ibm cn=user,ou=people,cn=aixdata`. The default value is `cn=admin`.
- c|--clusters**
Displays one or more clusters on which to run this command.
- d|--db2_password**
Specify the DB2[®] instance password.
- e|--encryption_seed**
Specify the encryption seed data for the new LDAP server configuration. This data is used to encrypt the key stash files and must contain at least 12 characters.
- h|-?|--help**
Requests help for this command.
- K|--ssl_password**
Specify the password for the Secure Sockets Layer (SSL) key.
- k|--key**
Specify the full path to the SSL key database, which must be in `.kdb` format. If you do not specify the path, an SSL key database is created automatically.
- p|--password**
Specify the text-only LDAP administrator password for the DN used to bind to the LDAP server. The password must match the password on the LDAP server for the specified DN.
- P|--password**
Specify the text-only LDAP bind password for the DN used to bind to the LDAP server. The password that you use must match the password on the LDAP server for the specified DN.
- S|--schema**
Displays the LDAP schema that is used to represent user and group entries in the LDAP server. When `rfc2307aix` is specified, the LDAP server is set up by using the RFC 2307 specifications and the auxiliary AIX schema. You must use this schema because of its interoperability through the RFC 2307 specifications and full AIX attribute support. You cannot edit the schema.

-t|--port

Specify the server port number that is used for communicating with the LDAP server. The default port number is 636.

-V|--version

Displays the current LDAP fileset's version. You must be using LDAP version 6.2.0.0, or later.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<server>

The host name that resolves to an LDAP server or the IP address of a server where a supported LDAP server is installed and running.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Configure a new LDAP server that is used by clusterA and is using the LDAP server that is running on myLDAP:

```
smcli sysmirror/mkldapsrv -c clusterA \  
    -a cn=admin -p myPassword \  
    -B "cn=aixdata,o=ibm" \  
    -k /keys/my_ssl_key -K my_ssl_password \  
    myLDAP
```

2. Configure an existing LDAP server that is used by clusterA and is using the LDAP server that is running on myLDAP:

```
smcli sysmirror/mkldapsrv -c clusterA \  
    -b cn=admin -P myPassword \  
    -B "cn=aixdata,o=ibm" \  
    -k /keys/my_ssl_key -K my_ssl_password \  
    myLDAP
```

moveresgrp command

Purpose

Use this command to move one or more resource groups from one node to another node.

Syntax

```
smcli sysmirror/moveresgrp [-h|-?|--help] [-v|--verbose]
smcli sysmirror/moveresgrp {-c|--cluster} <CLUSTER> \
                             {-n|--node} <NODE> \
                             [{-S|--state} {online|offline}] \
                             <RG>[,<RG#2>,...]
```

Command alias

mvrgr

Description

Moves a specified resource group from one node to another node. The resources that are managed by the specified resource group might be unavailable during the move process.

Flags

-c|--clusters

Displays one or more clusters on which to run this command.

-h|-?|--help

Requests help for this command.

-n|--node

The node that you want to move the resource group to.

-S|--state

The state of the resource group after it has been moved. You can specify if the resource group is online or offline. By default, the current state of the resource group is preserved.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<RG>

The label that you use to identify the resource group that you use for this command.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Move a resource group that is named RG1 to nodeB in a cluster named clusterA. After the move, the state of the resource groups remains the same.

```
smcli sysmirror/moveresgrp -c clusterA -n nodeB RG1
```

2. Move a resource group that is named RG2 to nodeC in a cluster named clusterA. After the move, the resource group is in an online state:

```
smcli sysmirror/mvrg -c clusterA -n nodeC -S online RG2
```

mknetwork command

Purpose

Creates a PowerHA SystemMirror network definition that conforms to the provided specifications.

Syntax

```
smcli sysmirror/mknetwork [-h|-?|--help] [-v|--verbose]
smcli sysmirror/mknetwork [-p|--public] \
    [{-t|--type} {Ether|Infiniband|XD_data|XD_ip}] \
    [{-N|--netmask} ###.###.###.###]\
    {-c|--cluster} <CLUSTER> \
    {-l|--label} <network_label>
smcli sysmirror/mknetwork [-p|--public] \
    [{-t|--type} {Ether|Infiniband|XD_data|XD_ip}] \
    {-P|--prefix} ### \
    {-c|--cluster} <CLUSTER> \
    {-l|--label} <network_label>
```

Command alias

mknw

Description

Creates a network definition within the specified cluster.

This network definition initially exists only on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of the network definition. To complete the creation, you must synchronize the cluster. This synchronization is done as a separate step so that you can make multiple changes to the cluster, and then verify those change. If you do not change the controlling node in the cluster, you can discard your changes.

Flags

-c|--clusters

Displays one or more clusters on which to run this command.

-h|-?|--help

Requests help for this command.

-N|--netmask

Specifies a netmask to use for the IPv4 persistent IP. The default netmask value is 255.255.255.0.

-P|--prefix

Specifies the prefix length to use for the IPv6 persistent IP.

-p|--public

Indicates that this network is to be public and it is not private.

-t|--type

Specifies the type of the network. Valid values are Ether, Infiniband, XD_data, XD_ip. The default value is Ether.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<network_name>

A label that you can use to identify the network.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Create an Ethernet network named myNetwork on cluster clusterA that has a netmask of 255.255.240.0:
smcli sysmirror/mknetwork -c clusterA -t Ether -N 255.255.240.0 myNetwork
2. Create an Ethernet network named myNetwork on cluster clusterA that has a prefix of 128:
smcli sysmirror/mknw -c clusterA -t Ether -P 128 myNetwork

mknode command

Purpose

Creates a PowerHA SystemMirror node definition that conforms to the provided specifications.

Syntax

```
smcli sysmirror/mknode [-h|-?|--help] [-v|--verbose]
smcli sysmirror/mknode [{-s|--site} <SITE>] \
                        [{-i|--comppath} <IP_ADDRESS_OR_HOSTNAME>] \
                        [-S|start_on_boot] \
                        [-B|broadcast_on_start] \
                        [-C|start_clinfo] \
                        [-V|verify_on_start] \
                        {-c|--cluster} <CLUSTER> \
                        <node_name>
```

Command alias

mknd

Description

Creates a node definition with in the specified cluster. The name given to the node can simply be the node's hostname. However, the name can also be something else, that is more descriptive. If the node name does not match the host name, the `--commpath` flag is required.

This new definition initially exists only on the controlling node of the specified cluster; none of the other nodes in the cluster are be aware of it yet. To complete the creation, you need to synchronize the cluster. This is done as a separate step so that you can make multiple changes to the cluster, and then verify those changes as a set before enabling them by synchronization. You can discard your changes at any time before the synchronization occurs, as long as you have not changed your cluster's controlling node.

Flags

-B|--broadcast_on_start

Indicates whether a terminal announcement must be made each time cluster services start.

-C|--start_clinfo

Indicates whether the clinfoES subsystem must be started when cluster services start.

-c|--cluster

The cluster on which to run this command.

-C|--confirm

Provides confirmation to preform this operation.

-h|-?|--help

Requests help for this command.

-i|--commpath

The communication path to take to the node, such as an IP address or host name.

-s|--site

A site to assign this node to.

-S|--start_on_boot

Indicates whether cluster services must be started when the host boots.

-v|--verbose

Requests maximum details in the displayed information.

-V|--verify_on_start

Indicates whether cluster verification must be performed when cluster services starts.

Parameters

<node_name>

The label of a node on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

1 The command has not been used in the intended manner.

2 A specified or required file was not found.

11 The command failed to accomplish its task.

12 A specified or a required resource was not found.

- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Create a simple node that is named `hostname` on `clusterA`:

```
smcli sysmirror/mknode -c clusterA hostname
```
2. Create a node with optional flags that is named `notTheHostname` on `clusterA`:

```
smcli sysmirror/mknode -c clusterA \  

-i 9.3.6.192 \  

-s Austin \  

-SBV \  


```

mkmethod command

Purpose

Creates a PowerHA SystemMirror method definition.

Syntax

```
smcli sysmirror/mkmethod [-h|-?|--help] [-v|--verbose]
smcli sysmirror/mkmethod {-c|--cluster} <CLUSTER> \  

{-t|--type} {snapshot|verify} \  

{-F|--file} <FILE> \  

[{-d|--description} "<DESCRIPTION>"] \  

<method_name>
smcli sysmirror/mkmethod {-c|--cluster} <CLUSTER> \  

{-t|--type} notify \  

{-C|--contact} <EMAIL> \  

{-e|--events} <EVENT>[,<EVENT#2>,...] \  

[{-n|--nodes} <NODE>[,<NODE#2>,...] ] \  

[{-F|--file} <FILE> ] \  

[{-d|--description} "<DESCRIPTION>"] \  

[{-r|--retry} ##] \  

[{-T|--timeout} ##] \  

<method_name>
```

Command alias

`mkmd`

Description

Creates a method definition within the specified cluster. A *method* in PowerHA SystemMirror is an object for managing an executable file, which can be a binary file or a script.

This method definition initially exists only on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of the method definition. To complete the creation of the method definition, you must synchronize the cluster. This synchronization is done as a separate step so that you can make multiple changes to the cluster and verify the changes. If you have not changed the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

- c|--clusters**
Displays one or more clusters on which to run this command.
- C|--contact**
Specifies one or more email addresses to send the notification to.
- d|--description**
A short one-line sentence that describes the object.
- e|--events**
One or more PowerHA SystemMirror event labels that, when triggered, results in the method being run.
- F|--file**
The executable file that is run by the method.
- h|-?|--help**
Requests help for this command.
- n|--nodes**
One or more node names that are authorized to send the remote notification for this method. Priority for nodes is from left to right. If the leftmost node is not active, the next node in the list attempts to send the remote notification.
- r|--retry**
The number of times to attempt to send the remote notification. The default value is 3.
- t|--timeout**
The number of seconds to wait for results from an attempt to send a remote notification. The default value is 45.
- t|--type**
The type of the method. You can set the type to notify, snapshot, or verify.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

- <method_name>**
A label that you can use to identify the method.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.

29 The current locale is not supported.

Examples

1. Create a snapshot method named `ssMethod` on cluster `clusterA` that runs a file named `/sm/methods/getDataScript.sh`:

```
smcli sysmirror/mkmethod -c clusterA -t snapshot \  
    -F /sm/methods/getDataScript.sh \  
    ssMethod
```

2. Create a verification method named `vrMethod` on cluster `clusterA` that runs a file named `/sm/methods/verifyScript.sh`:

```
smcli sysmirror/mkmethod -c clusterA -t verify \  
    -F /sm/methods/verifyScript.sh \  
    vrMethod
```

3. Create an event notification method named `tellMeMethod` on cluster `clusterA` that runs a file named `/sm/methods/tellMeScript.sh` each time a `node_down` or `node_down_local` event occurs. This example also shows how to send an email to `nobody@nowhere.com`:

```
smcli sysmirror/mkmethod -c clusterA -t verify \  
    -C nobody@nowhere.com \  
    -e node_down,node_down_local \  
    -F /sm/methods/tellMeScript.sh \  
    tellMeMethod
```

mkresgrp command

Purpose

Creates a PowerHA SystemMirror resource group definition within the specified cluster.

Syntax

```
smcli sysmirror/mkresgrp [-h|-?|--help] [-v|--verbose]  
smcli sysmirror/mkresgrp {-c|--cluster} <CLUSTER> \  
    {-n|--nodes} <NODE>[,<NODE#2>,...] \  
    [{-2|--secondary_nodes} <node>[,<node#2>,...]] \  
    [{-p|--site_policy} {ignore|primary|either|both} ] \  
    [{-s|--startup} {OHN|OFAN|OAN|OUDP}] \  
    [{-o|--fallover} {FNPN|FUDNP|BO}] \  
    [{-b|--fallback} {NFB|FBHPN}] \  
    [{-N|--node_priority} {default|mem|cpu| \  
        disk|least|most}] \  
    [{-y|--priority_script} <PATH_TO_FILE>] \  
    [{-T|--script_timeout} ###] \  
    [{-S|--service_ips} <SERVICE_IP>[,<SERVICE_IP#2>,...] ] \  
    [{-B|--fallback_timer} <FALLBACK_TIMER> ] \  
    [{-a|--app_controllers} <APPCTL>[,<APPCTL#2>,...] ] \  
    [{-t|--tapes} <TAPE>[,<TAPE#2>,...] ] \  
    [{-V|--volume_groups} <VG>[,<VG#2>,...] ] \  
    [-Y|--forced_varyon] \  
    [-i|--auto_import] \  
    [{-F|--file_systems} <FS>[,<FS#2>,...] ] \  
    [{-d|--disks} <DISK>[,<DISK#2>,...] ] \  
    [{-w|--wpar_name} <WPAR>] \  
    [{-e|--export_fs} <FS>[,<FS#2>,...] ] \  
    [{-E|--export_fs_v4} <FS>[,<FS#2>,...] ] \  
    [{-g|--stable_storage} <PATH_TO_STORAGE>] \  
    [{-W|--nfs_network} <NETWORK>] ] \  
    [{-m|--fs_mounts} <MOUNT_EXPR>[,<MOUNT_EXPR#2>,...] ] \  
    <resource_group_name>
```

Command alias

mkrg

Description

Creates a resource group definition on the specified controlling node.

This new definition initially exists only on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of it yet. To complete the creation process, the cluster must be synchronized by using the **synccluster** command.

This synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

-a|--app_controllers

One or more application controllers that will be started with this resource group.

-b|--fallback

The fallback policy for this resource group:

never fallback (NFB)

fallback to a higher priority node (FBHPN)

-c|--clusters

Displays one or more clusters on which to run this command.

-d|--disks

One or more PVIDs for disks that will be managed by this resource group.

-h|-?|--help

Requests help for this command.

-e|--export_fs

The mount points of the NFS v2 or NFS v3 file systems or directories, which are exported to all nodes in the resource chain when the resource group is initially acquired.

-E|--export_fs_v4

The mount points of the NFS v4 file systems or directories, which are exported to all nodes in the resource chain when the resource group is initially acquired.

-F|--file_systems

The mount points of the file systems, which should be mounted when the resource group is acquired. If file systems are not specified, and volume groups are specified for this resource group, by default all file systems that exist on the specified volume groups will be mounted when the resource group goes online. This default behavior of mounting all file systems is assumed if volume groups are explicitly specified as resources for this resource group. For any file system that is specified in this field, the volume group on which it resides will be mounted when the resource group is acquired.

-g|--stable_storage

NFS v4 needs stable storage to save state information. The path must belong to a file system, which is managed by the resource group. Do not export this path. This is to avoid accidental corruption of the NFS v4 state. Any change to the stable storage requires that the resource group is stopped and restarted.

-i|--auto_import

This value is **false** by default. If reset to **true**, the definition of any volume groups, which are managed by this resource group will be imported to any resource group nodes that do not already have it. The automatic import operation for volume groups is not supported for volume groups that contain data path devices (VPATHs).

When the automatic import operation is set to **true**, the final state of the volume group depends on the initial state of the volume group (varied on or varied off) and the state of the resource group to which the volume group is to be added (online or offline). When the automatic import operation is set to the default value of **false**, the final state of the volume group depends on the current state of the volume group and the state of the resource group to which the volume group is added.

-m|--fs_mounts

All nodes in the resource chain will attempt the NFS-mount operation on these file systems. Use the following syntax for specifying an NFS cross mount: `NFS_Mount_Point;Local_Fileystem`. For example, `/nfsmount/share1;/mnt/fs1`.

-N|--node_priority {default|mem|cpu|disk|least|most}

This field determines the priority order of nodes, which is used to determine the destination of a resource group during an event. This process causes the resource group to either move or be placed online.

The following options are possible:

- **default** - Uses the static priority order that is listed in the participating nodes list for this resource group (in left to right order of priority).
- **mem** - Ordered by the most memory available.
- **disk** - Ordered by the least disk activity.
- **cpu** - Ordered by most available CPU cycles.
- **least** - Ordered by lowest nonzero return value of the user-defined script.
- **most** - Ordered by highest return value of the user-defined script.

-n|--nodes

The participating nodes for this resource group in the order of priority (from left to right; highest priority to lowest priority).

-p|--site_policy

Indicate one of the following site policies:

- **ignore**
- **primary**
- **either**
- **both**

-2|--secondary_nodes

List the secondary nodes for this resource group. Each node in this list must be separated by a comma.

-o|--fallover

The fallover policies for this resource group follow:

- **FNPN** (fallover to the next priority node)
- **FUDNP** (fallover by using the dynamic node priority)
- **BO** (bring offline on error node only; concurrent resource group)

--S|--service_ips

One or more IP labels or addresses over which the services provided by resources of the resource group are typically provided, or over which clients connect to the applications of the resource group. These IP labels and IP addresses are kept highly available, and are moved with the resource group.

--B|--fallback_timer

One or more fallback timers names that are used by this resource group. This option is available only if the `--fallback` option has the **FBHPN** value.

-s|--startup

The startup policies for this resource group follow:

- OHN (on home node only)
- OFAN (on first available node)
- OAAN (on all available nodes; concurrent resource group)
- OUDP (online using distribution policy)

-T|--script_timeout

The maximum time, in seconds, during which the dynamic node priority script runs. This value must be less than the `config_too_long` time value. If the run time of the script exceeds this limit, the script stops running.

-t|--tapes

One or more tape resources that will be started with this resource group.

-W|--nfs_network

The preferred network used for the NFS mount operation of any specified file systems.

-v|--verbose

Requests maximum details in the displayed information.

-V|--volume_groups

One or more volume groups that contain raw logical volumes or raw volume groups that are varied on when the resource group is initially acquired.

-w|--wpar_name

If this attribute is set, all application controllers in the current resource group runs in the specified workload partition (WPAR). All service label and file system resources that are part of this resource group will be assigned to the specified WPAR. Changes to the WPAR name are not applied until the next time that the resource group is brought online. Do not specify a WPAR name if you do not want to enable this WPAR on the resource group.

-Y|--forced_varyon

If specified, PowerHA SystemMirror uses a forced varyon to bring the volume groups, which are managed by this resource group, online in the event that a normal varyon fails due to lack of quorum, and there is at least one complete copy of every logical volume that is available. This option is only meaningful for volume groups in which every logical volume is mirrored. You must use a super strict mirror allocation policy. This option will not be successful for other variations of logical volume configuration.

-y|--priority_script

The full path name of an executable file that is used for dynamic node priority.

Parameters

<resource_group_name>

A label by which the resource group might be referred to.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.

- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Create a resource group named rgA in cluster clusterA that incorporates nodes nodeA and nodeB:

```
smcli sysmirror/mkrg -c clusterA \
    -n nodeA,nodeB \
    rgA
```

2. Create a resource group named rgB in cluster clusterA that incorporates nodes nodeC and nodeD, and manages volume group myVG:

```
smcli sysmirror/mkrg -c clusterA \
    -n nodeC,nodeD \
    -V myVG \
    rgA
```

3. Create a concurrent (active on all nodes at the same time) resource group named rgConc in cluster clusterA that incorporates nodes nodeE and nodeF, and manages volume group myVG:

```
smcli sysmirror/mkrg -c clusterA \
    -n nodeE,nodeF \
    --startup OAAN \
    --failover B0 \
    --fallback NFB \
    rgConc
```

mksecurity command

Purpose

Configures the internode communication security within a PowerHA SystemMirror cluster.

Syntax

```
smcli sysmirror/mksecurity [-h|-?|--help] [-v|--verbose]
smcli sysmirror/mksecurity [{-l|--level} {Disable|Low|Med|High}] \
    [{-a|--algorithm} {DES|3DES|AES}] \
    [{-m|--mechanism} {OpenSSL|SelfSigned|SSH}] \
    [{-g|--grace_period} <SECONDS>] \
    [{-r|--refresh} <SECONDS>] \
    [{-k|--key} <PATH_TO_FILE>] \
    [{-c|--certificate} <PATH_TO_FILE>] \
    [-d|--auto_distrib] \
    <cluster_name>
```

Command alias

mksc

Description

Create or change security settings within the specified cluster.

Flags

-a|--algorithm

Indicates which symmetric algorithm to use in the security settings of the cluster. You can specify DES, 3DES, or AES.

- c|--certificate**
Indicates the path of the file containing the certificate that is used in the cluster security settings.
- d|--auto_distrib**
Indicates that the certificate is automatically distributed in the cluster security settings.
- g|--grace_period**
The grace period that is configured in the security settings of the cluster. The default value is 21600 seconds (6 hours).
- h|-?|--help**
Requests help for this command.
- k|--key**
Indicates the path of the file containing the key that is set in the cluster security settings.
- l|--level**
The level of the security to apply on the cluster. You can specify Disable, Low, Med, or High.
- m|--mechanism**
The mechanism to set in the cluster security settings. You can specify OpenSSL, SelfSigned, or SSH.
- r|--refresh**
The refresh time interval that is set in the security settings of the cluster. The default value is 86400 seconds (24 hours).
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

- <cluster_name>**
The label that you can use to identify the cluster.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Disable security on cluster clusterA:

```
smcli sysmirror/mksecurity -l disable clusterA
```
2. Establish high-level AES encryption on clusterA by using automatically created self-signed certificates:

```
smcli sysmirror/mksecurity -l high -a AES -m SelfSigned clusterA
```

mkserviceip command

Purpose

Creates a PowerHA SystemMirror service IP definition.

Syntax

```
smcli sysmirror/mkserviceip [-h|-?|--help] [-v|--verbose]
smcli sysmirror/mkserviceip {-c|--cluster} <CLUSTER> \
                             {-n|--network} <NETWORK> \
                             [{-N|--netmask} ###.###.###.###] \
                             [{-a|--hwaddr} <ADDR>] \
                             [{-s|--site} <SITE>] \
                             <service_ip>
smcli sysmirror/mkserviceip {-c|--cluster} <CLUSTER> \
                             {-n|--network} <NETWORK> \
                             [{-P|--prefix} ###] \
                             [{-a|--hwaddr} <ADDR>] \
                             [{-s|--site} <SITE>] \
                             <service_ip>
```

Command alias

mksi

Description

The service IP definition initially exists only on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of the service IP definition. To complete the creation of the service IP definition, you must synchronize the cluster.

This synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

-a|--hwaddr

Specifies the hardware address for the specified adapter. This address is used whenever another adapter assumes the IP address of the specified adapter. This flag eliminates the need for Address Resolution Protocol (ARP) cache updates, but increases the time for address swapping. This flag is not applicable for Internet Protocol version 6 (IPv6).

-c|--clusters

Displays one or more clusters on which to run this command.

-h|-?|--help

Requests help for this command.

-N|--netmask

A netmask that is used for IPv4 addresses and networks. The default netmask value is 255.255.255.0.

-P|--prefix

A prefix length that is used for IPv6 addresses and networks.

-s|--site

Indicate a service IP address that you can assign to a site.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<service_ip>

A resolvable hostname or IP address.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Create a service IP for the resolvable label mySIP on cluster clusterA that uses network net_ether_03:
smcli sysmirror/mkserviceip -c clusterA -n net_ether_03 mySIP
2. Create a service IP for IP address 10.14.217.199 on cluster clusterA that uses network net_ether_02:
smcli sysmirror/mkserviceip -c clusterA -n net_ether_02 10.14.217.199

mksite command

Purpose

Creates a PowerHA SystemMirror site definition that conforms to the provided specifications.

Syntax

```
smcli sysmirror/mksite [-h|-?|--help] [-v|--verbose]
smcli sysmirror/mksite {-c|--cluster} <CLUSTER> \
                       {-a|--multicast_address} <ADDRESS> \
                       [{-n|--nodes} <NODE>[,<NODE#2>,...] ] \
                       <site_name>
```

Description

Creates a site definition within the specified cluster.

It is important to understand that this removal only occurs on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of it yet. To complete the removal of the network, you must synchronize the cluster. This is done as a separate step so that you can make multiple changes to the cluster, and then verify those changes as a set before enabling them by using synchronization. You can discard your changes at any time before the synchronization occurs, if you have not changed your cluster's controlling node.

Flags

-c|--clusters

Displays one or more clusters on which to run this command.

-h|-?|--help

Requests help for this command.

-i|--site_ip

A multicast address that is used. The valid range is 224.0.0.0 - 239.255.255.255.

-n|--nodes

One or more nodes that you can add to the primary site.

-r|--repositories

Identifies the shared central repository database disks for the cluster. The repository can be identified as either a PVID or as an hdisk@agent.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<SITE>

The label of a site on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Create a site named Paris on clusterA:

```
smcli sysmirror/mksite -c clusterA \  
                        -a 240.111.222.121 \  
                        Paris
```

2. Create a site named siteB in cluster clusterA that incorporates nodes nodeC and nodeD, and uses a repository disk hdisk5 that has a multicast address of 226.1.2.3:

```
smcli sysmirror/mksite -c clusterA \  
                        -n nodeC,nodeD \  
                        -r hdisk5 \  
                        -i 226.1.2.3 \  
                        siteB
```

mksnapshot command

Purpose

Creates a PowerHA SystemMirror configuration snapshot that conforms to the provided specifications.

Syntax

```
smcli sysmirror/mksnapshot [-h|-?|--help] [-v|--verbose]
smcli sysmirror/mksnapshot {-c|--cluster} <CLUSTER> \
    {-d|--description} "<DESCRIPTION>" \
    [{-M|--methods} <METHOD>[,<METHOD#2>,...] ] \
    [-s|--save_logs] \
    <snapshot_name>
```

Command alias

mkss

Description

Creates a backup of the PowerHA SystemMirror configuration data for the specified cluster. Additional information might also be collected as part of the backup by using snapshot methods.

Flags

- c|--cluster**
A cluster on which this command runs.
- d|--description**
A brief description of the object.
- h|-?|--help**
Requests help for this command.
- M|--methods**
Enter one or more snapshot method names in the order in which they are to be run. Specifying **All** will result in all currently defined snapshot methods being run.
- s|--save_logs**
This option will collect cluster log files from all nodes in the cluster and save them in the snapshot. Saving log files can significantly increase the size of the snapshot.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

<snapshot_name>
A label by which the snapshot might be referred to.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.

- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. To create a snapshot named **snap1** on cluster `clusterA`, enter:
`smcli sysmirror/mksnapshot -c clusterA -d "My first snapshot." snap1`
2. To create a snapshot that uses a custom method named `method1`, enter:
`mcli sysmirror/mkss -c clusterA -d "My second snapshot." -M method1 snap2`

mktape command

Purpose

Creates a PowerHA SystemMirror tape drive definition.

Syntax

```
smcli sysmirror/mktape [-h|-?|--help] [-v|--verbose]
smcli sysmirror/mktape {-c|--cluster} <CLUSTER> \
                        {-D|--device} <DEVICE> \
                        [{-d|--description} "<DESCRIPTION>"] \
                        [{-s|--start} <PATH_TO_FILE>] \
                        [{-S|--sync_start}] \
                        [{-p|--stop} <PATH_TO_FILE>] \
                        [{-P|--sync_stop}] \
                        <tape_name>
```

Command alias

`mktp`

Description

The tape drive definition initially exists only on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of the tape drive definition. To create the tape drive definition, you must synchronize the cluster.

This synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

-c|--clusters

Displays one or more clusters on which to run this command.

-D|--device

The name of the tape special file. If the first character is not a `/`, `/dev/` is automatically added to the start of the file path.

-d|--description

A one-line sentence that describes the object.

- h|--help**
Requests help for this command.
- s|--start**
Specify an executable file that starts the tape resource.
- p|--stop**
Specify an executable file that stops the tape resource.
- S|--sync_start**
Indicates that the tape resource must be started synchronously.
- P|--sync_stop**
Indicates that the tape resource must be stopped synchronously.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

- <tape_name>**
A label that you can use to identify the tape device.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Create a tape object named myTape on cluster clusterA that manages the tape device that is installed at /dev/rmt0:

```
smcli sysmirror/mktape -c clusterA -D /dev/rmt0 myTape
```
2. Create a tape object named myTape on cluster clusterA that manages the tape device that is installed at /dev/rmt1. The tape uses a start script that is in the /scripts/start_tape1.sh directory and a stop script that is in the /scripts/stop_tape1.sh directory.

```
smcli sysmirror/mktape -c clusterA -D /dev/rmt1 \  

-s /scripts/start_tape1.sh \  

-P /scripts/stop_tape1.sh \  

myTape
```


mkuser command

Purpose

Creates an AIX user.

Syntax

```
smcli sysmirror/mkuser {-h|-?|--help} [-v|--verbose]
smcli sysmirror/mkuser {-c|--cluster} <CLUSTER> \
  {-R|--registry} {ldap|files} \
  [{-a|--authentication} {ldap|files} ] \
  [{-r|--resgrp} <resource_group>] \
  [{-i|--id} #####] \
  [{-O|--roles} <ROLE>[,<ROLE#2>,...]] \
  [-d|--admin_user] \
  [{-k|--keystore_access} {ldap|file}] \
  [{-A|--admin_keystore_access} {ldap|files}] \
  [{-m|--keystore_mode} {admin|guard}] \
  [{-C|--allow_mode_change}] \
  [{-K|--keystore_encryption} {RSA_1024|RSA_2048|
    RSA_4096}] \
  [{-F|--file_encryption} {AES_128_CBC|AES_128_ECB}] \
  [{-g|--groups} <group#1>[,<group#2>,...]] \
  [{-p|--primary} <group>] \
  [{-G|--admin_groups} <group#1>[,<group#2>,...]] \
  [-w|--no_switch_user] \
  [{-u|--su_groups} {ALL|<group#1>[,<group#2>,...]]] \
  [{-H|--home} <full_directory_path>] \
  [{-s|--shell} <defined_in_/etc/shells>] \
  [{-I|--inf} <user_information>] \
  [{-e|--expiration} <MMDDhhmmy>] \
  [-l|--locked] \
  [-L|--login] \
  [-o|--no_remote_login] \
  [{-S|--schedule} <range#1>[,<range#2>,...]] \
  [--max_failed_logins {#|0}] \
  [--allowed_ttys <tty#1>[,<tty#2>,...]] \
  [--days_to_warn {#|0}] \
  [--password_validation_methods <method#1>[,<md#2>,...]] \
  [--password_filters <filter#1>[,<filter#2>,...]] \
  [--min_passwords <##>] \
  [--reuse_time <##>] \
  [--lockout_delay <##>] \
  [--max_password_age {0..52}] \
  [--min_password_age {0..52}] \
  [--min_password_length {0..8}] \
  [--min_password_alphas {0..8}] \
  [--min_password_others {0..8}] \
  [--max_password_repeated_chars {0..52}] \
  [--min_password_different {0..8}] \
  [--password_registry {files|ldap|NIS|DCE}] \
  [{-U|--umask} #####] \
  [--audit_classes <class#1>[,<class#2>,...]] \
  [{-t|--trusted_path} {nosak|on|notsh|always}] \
  [--primary_auth {SYSTEM|NONE|<method>}] \
  [--secondary_auth {NONE|SYSTEM|<method>}] \
  <user_name>
```

Command alias

mkur

Description

Creates an AIX user on all the nodes in the cluster, or on all the nodes in a specified resource group.

Flags

-A|--admin_keystore_access

Specify where the Encrypted File System (EFS) admin keystore file is located when a user opens their keystore. The keystore allows files to be used in an EFS. The selection of a file creates a keystore file that is associated with the specified user. The default value indicates that keystore attributes are stored in LDAP.

-a|--authentication

Specify the mechanism that is used to authenticate users. The default value is `ldap_auth`. If `ldap_auth` is specified, you must use SSL to protect clear text passwords from displaying. If `unix_auth` is specified, user passwords are retrieved from the LDAP server and comparison is performed locally.

--allowed_ttys

Specifies the list of terminals that can access the specified user account. When a user tries to access the account, the system attempts to match the terminal from which the access request is made with a terminal listed in this field. The system works through the list of TTYs in the order that is specified in this field. The system grants access to the account and to the first TTY that it matches.

If the system cannot find a match, the user cannot log in to the account from that terminal. To enter a list of valid TTYs for this user account, type in the full path names to each terminal, `dev/tty0`, separating each path name with a comma. To exclude a specific TTY from being able to access this user account, precede it with an exclamation point. The special keyword `ALL` is the default value, and it indicates that the user can log in from any terminal.

--audit_classes

Specify the audit classes for the user by using a comma-separated list of classes, or a value of `ALL` to indicate all audit classes.

-C|--allow_mode_change

Specify the allow mode change.

-c|--clusters

Displays one or more clusters on which to run this command.

-d|--admin_user

Indicates that the user is an administrator. Only the root user can change the attributes of users that are defined as administrators.

--days_to_warn

Specifies the number of days before the expiration of the user's password when a warning message is issued. The message is displayed each time that the user logs in during this warning period. The message also displays the date when the password expires. The default value is 0.

-e|--expiration

Specify the expiration date of the account, after which the account is locked. The parameter is a 10-character string in the `MMDDhhmmyy` format, where `MM` = month, `DD` = day, `hh` = hour, `mm` = minute, and `yy` = last 2 digits of the years 1939 - 2038. All characters are a numeric value. If you specify 0, the account does not expire. The default value is 0.

-F|--file_encryption

Select the encryption algorithm that is used for encrypting the files for the specified user when they are created in an EFS. The default value is `AES_128_CBC`.

-G|--admin_groups

Specify the groups that the administrators can modify by using a comma-separated list for the group names.

-g|--groups

Specify the names of the groups that the specified user belongs to. *Groups* are collections of users that can share access authority to protected resources. Groups can be formed for users who access the same applications or hardware resources, perform similar tasks, or have similar needs for information. A user can be a member in up to 64 groups.

-H|--home

Specify the full path to the home directory for the specified user.

-h|--help

Requests help for this command.

-I|--info

Specify general information about the user. This flag is a string without any embedded colon (:) characters and cannot end with the number sign (#) or exclamation point (!) characters.

-i|--id

Specify a unique decimal integer string to associate with this user account on the system. If you do not specify a user ID, a user ID is automatically specified. The specified user ID is created on all nodes in the cluster or in the LDAP server.

Note: If you do not specify this value, AIX can assign different user IDs on each node.

-K|--keystore_encryption

Select the algorithm that is used to generate the private key, which is used during the keystore creation process. This key is used to protect the files for the specified user within an EFS. The default encryption value is `RSA_1024`.

-k|--keystore_access

Specify the database type of the keystore. This database is where the keystore attributes are stored. The keystore can be stored remotely on an LDAP server, or on a locally managed shared file system.

-L|--login

Specify whether the user can log in to the system by using the `login` command.

-l|--locked

Specify whether you want the user account to be initially locked, which prevents the user from logging into the system.

--lockout_delay

The number of weeks after the user password expires during which the user can still change the password. If this time period passes without a password change, the user cannot log in to their account until an administrator resets the password. If a value of 0 is specified, the user cannot log in at the time the password expires. If a value of -1 is specified, features of this flag are disabled. This flag is ignored if the `max_password_age` flag value is set to 0.

--max_password_age

Specify the maximum age in weeks for the user password. When the password reaches the specified age, the system requires it to be changed before the user can log in again. If you specify 0, the function is disabled and the password does not expire. You can specify a numeric value 0 - 52. The default value is 0.

--max_password_repeated_chars

Specify the maximum number of times that a character can be repeated in the user's password. If you specify 8, any number of characters can be repeated. You can specify a number in the range 0 - 8.

--min_password_age

Specify the minimum age in weeks for the user password before it can be changed. If you specify 0, the password can be changed at any time. You can specify a number in the range 0 - 52. The default value is 0.

--min_password_alphas

Specify the minimum number of alphabetical characters that must be included in the user password. If you specify 0, no minimum number of alphabetical characters is required. You can specify a number in the range 0 - 8. The default value is 0.

--min_password_different

The minimum number of characters that are required in the new password that were not in the old password. The value is a decimal integer string. If you specify 0, a minimum number of different characters is not required. You can specify a number in the range 0 - 8. The default value is 0.

--min_password_length

The minimum number of characters that the user password must contain. If you specify 0, there is no minimum length. You can specify a number 0 - 8. The default value is 0.

Note: The minimum length of a password is determined by the **min_password_length** parameter or the **min_password_alphas** parameter and the **min_password_others** parameter, whichever is greater. The **min_password_alphas** parameter and the **min_password_others** parameter must never be greater than 8. If the **min_password_alphas** parameter and the **min_password_others** parameter is greater than 8, the effective value for **min_password_others** parameter is reduced to 8 minus the **min_password_alphas** parameter.

--min_password_others

The minimum number of characters other than alphabetical characters that must be included in the user password. If you specify 0, a minimum number of other characters is not required. You can specify a number in the range 0 - 8. The default value is 0.

--min_passwords

The number of previous passwords that the user cannot reuse. This value might depend on the value of the **--reuse_time** flag. If you specify 0, any previous password can be reused provided that the specified number of weeks before the password reuse has occurred.

-o|--roles

Specify a role for the user from the list of available PowerHA SystemMirror roles. *Roles* define the tasks that a user can access and perform. Use this flag if the registry type is LDAP.

-o|--no_remote_login

Indicates that the user cannot log in from a remote location with the **telnet** or **rlogin** commands. By default, the users can log in remotely.

-p|--primary

Specify the user's primary group. If you do not specify a group, a default group is automatically selected.

--password_filters

Specify a list of dictionary files, in absolute path format, which contains words that cannot be used as passwords. The list must be a comma-separated list.

--password_validation_methods

List of administrator-supplied methods for checking the new password during a password change. The value is a comma-separated list of program names, which must be specified by using absolute path names, or a path relative to the `/usr/lib` directory. If the password does not meet the requirements of all the specified methods, the password cannot be changed.

--primary_auth

Lists the primary methods for authenticating the user. The parameters might be a comma-separated list of methods. The **method** parameter is the name of the authentication method. Valid authentication methods are defined in the `/etc/security/login.cfg` file. The SYSTEM method and local password authentication are the default parameters. If you specify the NONE method, primary authentication check does not occur.

--secondary_auth

Lists the secondary methods for authenticating the user. The parameters might be a comma-separated list of methods. The **method** parameter is the name of the authentication method. If you do not specify a parameter, the default of NONE is used. The NONE method indicates that a secondary authentication does not occur. Valid authentication methods are defined in the `/etc/security/login.cfg` file.

-R|--registry

Indicates where authentication information and credentials are to be stored.

-r|--resgrp

Specify a resource group to verify that the user is created on all the nodes that are defined within the specified resource group. This parameter can only be specified when the registry value is not set to LDAP.

--reuse_time

The number of weeks that must pass before a user is able to reuse a password after it has been selected as their current password. The interpretation of this parameter can depend on the value specified for the **--min_passwords** flag. If 0 is specified, any previous password can be reused if the minimum number of passwords before the reuse value has been reached.

-S|--schedule

Specify the days and times that the user is allowed to access the account by using a comma-separated list of items in the `MMdd[-MMdd]:hhmm-hhmm` format or the `w[w]:hhmm-hhmm` format, where MM = month, DD = day, hh = hour, mm = minute, and w = week. Entries that start with the exclamation point (!) are denied access to the system and are called DENY entries. Entries that do not start with the exclamation point (!) are allowed access to they system and are called ACCESS entries. The exclamation point (!) prefix applies to single entries and must be used at the start of each entry. You can use 200 entries per user.

-s|--shell

Specify the full path to the program that is run for the user at the start of a session.

-t|--trusted_path

Specifies a variable that defines access to the trusted path for the user. The system uses this variable when the user calls the trusted shell or a trusted process, or enters the secure attention key (SAK) sequence. You can specify the following values:

nosak The SAK sequence is disabled for all processes that are run by the user. Use this option if this user transfers binary data that might contain the SAK sequence.

notsh User cannot start the trusted shell on a trusted path and is logged off the system if the user enters the SAK sequence during the current session.

always

User can only execute trusted processes and the initial start program for the user is a trusted shell or another trusted process.

on User has standard trusted path characteristics and can start a trusted shell with the SAK sequence.

-U|--umask

Determines the default permissions on all files and directories that the specified user creates. This value, along with the permissions of the creating process, determine the permissions on a file or directory when it is created. The default value is 022.

-u|--su_groups

Specifies the groups that can switch to the this user account by using the **su** command. You might want groups, such as a group with administrative privileges, to access this user account to update the system configuration for the user. The special keyword ALL can be used to indicate that all groups are allowed use the **su** command to switch to this user, or an exclamation point (!) can be appended to a group name to exclude that specific group.

-v|--verbose
Requests maximum details in the displayed information.

-w|--no_switch_user
Indicates that no other users can switch to this user account by using the **su** command. By default, other users can switch to this user account.

Parameters

<user_name>
A label that you can use to identify the user.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Create a user named `hsimpson` on a cluster called `clusterA` that is stored in the local file system:
`smcli sysmirror/mkuser -c clusterA -R files hsimpson`
2. Create a user named `msimpson` on a cluster called `clusterA` that is stored on LDAP server in `clusterA`:
`smcli sysmirror/mkuser -c clusterA -R ldap msimpson`

movresgrp command

Purpose

Moves one or more PowerHA SystemMirror resource groups to a new node within the specified cluster.

Syntax

```
smcli sysmirror/moveresgrp [-h|-?|--help] [-v|--verbose]
smcli sysmirror/moveresgrp {-c|--cluster}<CLUSTER> \
                             {-n|--node} <NODE> \
                             {-S|--state} {online|offline} \
                             <RG>[,<RG#2>,...]
```

Command alias

`mvrg`

Description

Moves the specified resource groups from one node to another. The resources managed by the specified resource groups might be unavailable during the move process.

Flags

-c|--cluster

The cluster on which to run this command.

-h|-?|--help

Requests help for this command.

-n|--node

A node on which to perform this operation.

-S|--state

The desired state of the resource group after the move, online or offline operation. By default, an attempt is made to preserve the current state of the resource group.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<RG>

The label of a resource group on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Move a resource group named db2grp to node nodeD:

```
smcli sysmirror/moveresgrp -c clusterA \  
                             -n nodeD \  
                             db2grp
```

2. Move a resource group named db2grp to node nodeB, bringing it online on its new node:

```
smcli sysmirror/moveresgrp --cluster clusterA \  
                             --node nodeD \  
                             --state online \  
                             db2grp
```

rfcluster command

Purpose

Refreshes the information about one or more clusters in the IBM Systems Director database.

Syntax

```
smcli sysmirror/rfcluster {-h|-?|--help} \
                          [-v|--verbose]
smcli sysmirror/rfcluster <CLUSTER>[,<CLUSTER#2>,...]
smcli sysmirror/rfcluster [-a|--all]
```

Command alias

rfcl

Description

Refreshes the information in the IBM Systems Director database. This operation can be helpful if the IBM Systems Director database is out of sync with one or more clusters. Refreshing can also be helpful to overcome the delay that can occur between changes taking place in a cluster, and those changes showing up in the database. Refreshing cluster data should not be necessary under normal circumstances, but if it must be done, be aware that it can take a long time.

Flags

-a|--all

Indicates that this operation must be performed on all instances of this class.

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<CLUSTER>

The label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Refresh the data for clusterA:
`smcli sysmirror/rfcluster clusterA`
2. Refresh the data for all clusters that are currently managed by IBM Systems Director:

Note: Running the following command can take a long time to complete.

```
smcli sysmirror/rfcluster --all
```

rmappctl command

Purpose

Removes one or more IBM SystemMirror PowerHA application controllers from the specified cluster.

Syntax

```
smcli sysmirror/rmappctl [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmappctl {-c|--cluster}<CLUSTER> \
    [-C|--confirm] \
    <APPCTL>[,<APPCTL#2>,...]
smcli sysmirror/rmappctl {-c|--cluster}<CLUSTER> \
    [-C|--confirm] \
    {-a|--all}
```

Command alias

rmac

Description

Completely removes one or more networks from the specified cluster. If no confirmation of the removal is given by using the `--confirm` flag, the removal does not occur. You should create a snapshot of the affected cluster before running this command.

It is important to understand that this removal only occurs on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of it yet. To complete the removal of the network, you must synchronize the cluster. This is done as a separate step so that you can make multiple changes to the cluster, and then verify those changes as a set before enabling them by using synchronization. You can discard your changes at any time before the synchronization occurs, if you have not changed your cluster's controlling node.

Flags

- a|--all**
Indicates that this operation should be performed on all instances of this class.
- c|--cluster**
The cluster on which to run this command.
- C|--confirm**
Provides confirmation to perform the operation.
- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

<APPCTL>

The label of an application controller to perform this operation on.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Remove an application controller named myApp from a cluster named clusterA:
smcli sysmirror/rmappctl --confirm -c clusterA myApp
2. Remove all application controllers that are defined within cluster clusterA:

Note: You must run this command only after you have planned for the removal. If you remove the fallback timer, your environment is affected.

```
smcli sysmirror/rmappctl -C -c clusterA --all
```

rmappmon command

Purpose

Removes one or more PowerHA SystemMirror application monitors from the specified cluster.

Syntax

```
smcli sysmirror/rmappmon [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmappmon {-c|--cluster}<CLUSTER> \
[-C|--confirm] \
<APPMON>[,<APPMON#2>,...]
smcli sysmirror/rmappmon {-c|--cluster}<CLUSTER> \
[-C|--confirm] \
{-a|--all}
```

Command alias

rmam

Description

Completely removes one or more networks from the specified cluster. If no confirmation of the removal is given by using the `--confirm` flag, the removal does not occur. You should create a snapshot of the affected cluster before running this command.

It is important to understand that this removal only occurs on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of it yet. To complete the removal of the network, you must synchronize the cluster. This is done as a separate step so that you can make multiple changes to the cluster, and then verify those changes as a set before enabling them by using synchronization. You can discard your changes at any time before the synchronization occurs, if you have not changed your cluster's controlling node.

Flags

-a|--all

Indicates that this operation should be performed on all instances of this class.

-c|--cluster

The cluster on which to run this command.

-C|--confirm

Provides confirmation to perform the operation.

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<APPMON>

The label of an application monitor to perform this operation on.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Remove an application monitor named `procmon` from a cluster named `clusterA`:
`smcli sysmirror/rmappmon --confirm -c clusterA procmon`

2. Remove all application monitors defined within cluster `clusterA`:

Note: You must run this command only after you have planned for the removal. If you remove the fallback timer, your environment is affected.

```
smcli sysmirror/rmappmon -C -c clusterA --all
```

rmcluster command

Purpose

Removes one or more PowerHA SystemMirror clusters completely from all their member nodes.

Syntax

```
smcli sysmirror/rmcluster [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmcluster [-C|--confirm] \
    <CLUSTER>[,<CLUSTER#2>,...]
smcli sysmirror/rmcluster [-C|--confirm] \
    {-a|--all}
```

Command alias

`rmcl`

Description

Completely removes one or more networks from the specified cluster. If no confirmation of the removal is given by using the `--confirm` flag, the removal does not occur. You should create a snapshot of the affected cluster before running this command.

Flags

`-a|--all`

Indicates that this operation should be performed on all instances of this class.

`-C|--confirm`

Provides confirmation to perform this operation.

`-h|-?|--help`

Requests help for this command.

`-v|--verbose`

Requests maximum details in the displayed information.

Parameters

`<CLUSTER>`

The label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

1 The command has not been used in the intended manner.

2 A specified or required file was not found.

11 The command failed to accomplish its task.

12 A specified or a required resource was not found.

- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Remove a cluster named `clusterA`:

```
smcli sysmirror/rmcluster --confirm clusterA
```
2. Remove all PowerHA SystemMirror clusters that are currently managed by IBM Systems Director:

Note: Only run this command after you have planned for the removal and are aware of how the removal will affect your environment.

```
smcli sysmirror/rmcluster -c --all
```

rmdependency command

Purpose

Eliminates one or more resource group dependencies from the specified cluster.

Syntax

```
smcli sysmirror/rmdependency [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmdependency {-c|--cluster} <CLUSTER> \
[-C|--confirm] \
[{-t|--type} {PARENT_CHILD|STOP_AFTER| \
START_AFTER|SAME_NODE| \
SAME_SITE|DIFFERENT_NODES}] \
<DEPENDENCY>[ <DEPENDENCY#2> ...]
smcli sysmirror/rmdependency {-c|--cluster} <CLUSTER> \
[-C|--confirm] \
[{-t|--type} {PARENT_CHILD|STOP_AFTER| \
START_AFTER|SAME_NODE| \
SAME_SITE|DIFFERENT_NODES}] \
{-a|--all}
```

Command alias

`rmdp`

Description

If you do not use the **--confirm** flag, a warning message is displayed asking you to use the **--confirm** flag.

This synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

-a|--all

Indicates that this operation is performed on all instances of the class.

-c|--clusters

Displays one or more clusters on which to run this command.

-C|--confirm
Displays information that the operation completed successfully.

-h|-?|--help
Requests help for this command.

-t|--type
Specify one of the following types:

- DIFFERENT_NODES
- PARENT_CHILD
- SAME_NODE
- SAME_SITE
- START_AFTER
- STOP_AFTER

-v|--verbose
Requests maximum details in the displayed information.

Parameters

<DEPENDENCY>
The label that you can use to identify a dependency.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Remove a parent dependency and a child dependency named RG1+RG2 from clusterA:
`smcli sysmirror/rmdependency --confirm -c clusterA RG1+RG2`
2. Remove all dependencies that are defined within clusterA:

Note: Before you run this command, you must plan for the removal of the cluster. If you remove the cluster, your environment is affected.

```
smcli sysmirror/rmdp -C -c clusterA --all
```

rmefsKeystore command

Purpose

Eliminates the Encrypted File System (EFS) keystore from a cluster.

Syntax

```
smcli sysmirror/rmefsKeystore [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmefsKeystore [-C|confirm] <CLUSTER>
```

Command alias

rmef

Description

Deletes the EFS keystore from the cluster. PowerHA SystemMirror cannot manage the EFS keystore after it is deleted.

If you do not use the **--confirm** flag, a warning message is displayed asking you to use the **--confirm** flag.

Flags

-C|--confirm

Provides confirmation to perform the operation.

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<CLUSTER>

Specifies the label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Example

The following example removes the EFS keystore from clusterA:

```
smcli sysmirror/rmef --confirm clusterA
```

rmfallbacktimer command

Purpose

Removes one or more fallback timers from a cluster.

Syntax

```
smcli sysmirror/rmfallbacktimer [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmfallbacktimer {-c|--cluster} <CLUSTER> \
    [-C|--confirm] \
    <FALLBACK_TIMER>[,<FALLBACK_TIMER#2>,...]
smcli sysmirror/rmfallbacktimer {-c|--cluster} <CLUSTER> \
    [-C|--confirm] \
    {-a|--all}
```

Command alias

rmft

Description

If you do not use the **--confirm** flag, a warning message is displayed asking you to use the **--confirm** flag.

You must create a snapshot of the affected cluster for backup purposes before you run this command.

The removal of the fallback timer occurs on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of the removal of the fallback timer. To complete the removal of the fallback timer, you must synchronize the cluster.

This synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

-a|--all

Indicates that this operation must be performed on all instances of this class.

-c|--clusters

Displays one or more clusters on which to run this command.

-C|--confirm

Provides confirmation to perform the operation.

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<FALLBACK_TIMER>

The label that you can use to identify the fallback timer.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Remove a fallback timer named `hourlyTimer` from `clusterA`:

```
<p>smcli sysmirror/rmft --confirm -c clusterA hourlyTimer</p>
```
2. Remove all fallback timers that are defined within `clusterA`:

Note: You must run this command only after you have planned for the removal. If you remove the fallback timer, your environment is affected.

```
smcli sysmirror/rmft -C -c clusterA --all
```

rmfilecollection command

Purpose

Removes one or more PowerHA SystemMirror file collections from the specified cluster.

Syntax

```
smcli sysmirror/rmfilecollection [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmfilecollection {-c|--cluster}<CLUSTER> \
    [-C|--confirm] \
    <FC>[,<FC#2>,...]
smcli sysmirror/rmfilecollection {-c|--cluster}<CLUSTER> \
    [-C|--confirm] \
    {-a|--all}
```

Command alias

`rmfc`

Description

Completely removes one or more file collections from the specified cluster. None of the files or directories managed by the deleted file collections are removed. You should create a snapshot of the affected cluster is recommended before running this command.

It is important to understand that this removal only occurs on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of it yet. To complete the removal of the

network, you must synchronize the cluster. This is done as a separate step so that you can make multiple changes to the cluster, and then verify those changes as a set before enabling them by using synchronization. You can discard your changes at any time before the synchronization occurs, if you have not changed your cluster's controlling node.

If you do not use the **--confirm** flag, a warning message is displayed asking you to use the **--confirm** flag.

Flags

-a|--all

Indicates that this operation is performed on all instances of this class.

-c|--cluster

The cluster on which to run this command.

-C|--confirm

Provides confirmation to perform the operation.

-h|-?|--help

Requests help for this command.

v|--verbose

Requests maximum details in the displayed information.

Parameters

<FC>

The label of a file collection on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Remove a file collection named `myfc` from a cluster named `clusterA`:
`smcli sysmirror/rmfilecollection --confirm -c clusterA myfc`
2. Remove all file collections that are defined within cluster `clusterA`:

Note: You must run this command only after you have planned for the removal. If you remove the fallback timer, your environment is affected.

```
smcli sysmirror/rmfilecollection -C -c clusterA --all
```

rmgroup command

Purpose

Removes one or more AIX user groups from a cluster.

Syntax

```
smcli sysmirror/rmgroup [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmgroup {-c|--cluster} <CLUSTER> \
                        {-R|--registry} <ldap|files> \
                        [{-r|--resgrp} <resource_group>] \
                        [-C|--confirm] \
                        <group_name>[,<group_name#2>,...]
```

Command alias

rmgr

Description

If confirmation of the removal is not displayed when you use the **--confirm** flag, the AIX user group is not removed. You should create a snapshot of the affected cluster for backup purposes before you run this command.

Flags

-c|--clusters

Displays one or more clusters on which to run this command.

-C|--confirm

Provides confirmation to perform the operation.

-h|-?|--help

Requests help for this command.

-R|--registry

You can use one the following options to indicate from where to remove the user group definition from:

ldap Removes all information and credentials that are stored in an LDAP server.

files Removes all information and credentials that are stored locally in the cluster.

-r|--resgrp

Specify a resource group to ensure that the specified groups are removed from all the nodes that are defined within the resource group.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<group_name>

A label that you can use to identify the AIX user group.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

1 The command has not been used in the intended manner.

- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Remove a local group named myGroup from nodes that manage resource group RG1 in clusterA:

```
smcli sysmirror/rmggroup --confirm -R files -r RG1 -c clusterA myGroup
```
2. Remove an LDAP group named myGroup from clusterA:

```
smcli sysmirror/rmggroup --confirm -R ldap -c clusterA myGroup
```

rminterface command

Purpose

Removes one or more PowerHA SystemMirror interface definitions from the specified cluster.

Syntax

```
smcli sysmirror/rminterface [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rminterface {-c|--cluster}<CLUSTER> \
                             [-C|--confirm] \
                             <INTERFACE>[,<INTERFACE#2>,...]
smcli sysmirror/rminterface {-c|--cluster}<CLUSTER> \
                             [-C|--confirm] \
                             {-a|--all}
```

Command alias

rmif

Description

Completely removes one or more networks from the specified cluster. If no confirmation of the removal is given by using the `--confirm` flag, the removal does not occur. You should create a snapshot of the affected cluster before running this command.

It is important to understand that this removal only occurs on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of it yet. To complete the removal of the network, you must synchronize the cluster. This is done as a separate step so that you can make multiple changes to the cluster, and then verify those changes as a set before enabling them by using synchronization. You can discard your changes at any time before the synchronization occurs, if you have not changed your cluster's controlling node.

Flags

- a|--all**
Indicates that this operation should be performed on all instances of this class.
- c|--cluster**
The cluster on which to run this command.

- C|--confirm**
Provides confirmation to perform the operation.
- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

<INTERFACE>

The label of an interface on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Remove an interface named `myint` from a cluster named `clusterA`:
`smcli sysmirror/rminterface --confirm -c clusterA myint`
2. Remove all interfaces defined within cluster `clusterA`:

Note: You must run this command only after you have planned for the removal. If you remove the fallback timer, your environment is affected.

```
smcli sysmirror/rminterface -C -c clusterA --all
```

rmldapcl command

Purpose

Removes the LDAP client configuration from all the nodes of the cluster.

Syntax

```
smcli sysmirror/rmldapcl [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmldapcl {-C|--confirm} <CLUSTER>
```

Command alias

`rmcl`

Description

If you do not use the **--confirm** flag, a warning message is displayed asking you to use the **--confirm** flag.

Flags

- C|--confirm**
Provides confirmation to perform the operation.
- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

- <CLUSTER>**
Specifies the label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Example

The following example removes an LDAP client configuration from clusterA:

```
smcli sysmirror/rmldapcl --confirm clusterA
```

rmldapsrv command

Purpose

Removes the LDAP server configuration from the cluster.

Syntax

```
smcli sysmirror/rmldapsrv [-h|-?|--help] [-v|--verbose]  
smcli sysmirror/rmldapsrv {-C|--confirm} <CLUSTER>
```

Command alias

rmsr

Description

If you do not use the **--confirm** flag, a warning message is displayed asking you to use the **--confirm** flag.

Flags

-C|--confirm

Provides confirmation to perform the operation.

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<CLUSTER>

Specifies the label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Example

The following example removes the LDAP server configuration from clusterA:

```
smcli sysmirror/rmldapsrv --confirm clusterA
```

rmnetwork command

Purpose

Removes one or more PowerHA SystemMirror networks from the specified cluster.

Syntax

```
smcli sysmirror/rmnetwork [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmnetwork --cluster <CLUSTER> \
    [-C|--confirm] \
    <NETWORK>[,<NETWORK#2>,...]
smcli sysmirror/rmnetwork --cluster <CLUSTER> \
    [-C|--confirm] \
    {-a|--all}
```

Command alias

rmnw

Description

Completely removes one or more networks from the specified cluster. If no confirmation of the removal is given by using the `--confirm` flag, the removal does not occur. You should create a snapshot of the affected cluster before running this command.

It is important to understand that this removal only occurs on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of it yet. To complete the removal of the network, you must synchronize the cluster. This is done as a separate step so that you can make multiple changes to the cluster, and then verify those changes as a set before enabling them by using synchronization. You can discard your changes at any time before the synchronization occurs, if you have not changed your cluster's controlling node.

Flags

-a|--all

Indicates that this operation should be performed on all instances of this class.

-c|--cluster

The cluster on which to run this command.

-C|--confirm

Provides confirmation to perform the operation.

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<NETWORK>

The label of a network on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.

- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Remove a network named myNetwork from a cluster named clusterA:
`smcli sysmirror/rmnetwork --confirm -c clusterA myNetwork`
2. Remove all networks that are defined within cluster clusterA:

Note: Only run this command after you have planned for the removal and are aware how the removal will affect your environment.

```
smcli sysmirror/rmnetwork -c -c clusterA --all
```

rmmethod command

Purpose

Removes one or more methods from the cluster.

Syntax

```
smcli sysmirror/rmmethod [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmmethod {-c|--cluster} <CLUSTER> \
                        [-C|--confirm] \
                        [{-t|--type} {notify|snapshot|verify}] \
                        <METHOD>[,<METHOD#2>,...]
smcli sysmirror/rmmethod {-c|--cluster} <CLUSTER> \
                        [-C|--confirm] \
                        [{-t|--type} {notify|snapshot|verify}] \
                        {-a|--all}
```

Command alias

rmmd

Description

If you do not use the **--confirm** flag, a warning message is displayed asking you to use the **--confirm** flag.

Before you run this command, you should create a snapshot of the affected cluster for backup purposes.

The removal of the method occurs on the controlling node of the specified cluster. None of the other nodes in the cluster are aware the removal of the method. To complete the removal of the method, you must synchronize the cluster.

This synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

-a|--all

Indicates that this operation must be performed on all instances of this class.

- c|--clusters**
Displays one or more clusters on which to run this command.
- C|--confirm**
Provides confirmation to perform the operation.
- h|-?|--help**
Requests help for this command.
- t|--type**
The type of the method. You can set the type to notify, snapshot, or verify.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

- <METHOD>**
The label that you can use to identify a method.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Remove a snapshot method named mySSMethod from clusterA:
`smcli sysmirror/rmmethod --confirm -c clusterA -t snapshot mySSMethod`
2. Remove all snapshot methods that are defined within clusterA:

Note: You must run this command only after you have planned for the removal. If you remove the fallback timer, your environment is affected.

```
mcli sysmirror/rmmd -C -c clusterA --all
```

rmnode command

Purpose

Removes one or more PowerHA SystemMirror nodes from the specified cluster.

Syntax

```
smcli sysmirror/rmnode [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmnode {-c|--cluster} <CLUSTER> \
                        [-C|--confirm] \
                        <NODE>[,<NODE#2>,...]
smcli sysmirror/rmnode --cluster <CLUSTER> \
                        [-C|--confirm] \
                        {-a|--all}
```

Command alias

rmnd

Description

Completely removes one or more nodes from the specified cluster. If no confirmation of the removal is given by the `--confirm` flag, the removal do not occur. You should create a snapshot of the affected cluster before running this command.

The removal only occurs on the controlling node of the specified cluster; none of the other nodes in the cluster are aware of it yet. To complete the removal, you need to synchronize the cluster. This is done as a separate step so that you can make multiple changes to the cluster, then verify those changes as a set before enabling them by using synchronization. You can discard your changes at any time before the synchronization occurs, as long as you have not changed your cluster's controlling node.

Flags

- a|--all**
Indicates that this operation should be performed on all instances of this class.
- c|--cluster**
The cluster on which to run this command.
- C|--confirm**
Provides confirmation to perform the operation.
- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

- <NODE>**
The label of a node on which to perform this operation.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.

- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Remove a node named `myNode` from a cluster named `clusterA`:
`smcli sysmirror/rmnode --confirm -c clusterA myNode`
2. Remove all nodes defined within cluster `clusterA`:

Note: Only run this command after you have planned for the removal and are aware how the removal will affect your environment.

```
smcli sysmirror/rmnode -c -c clusterA --all
```

rmpersistentip command

Purpose

Removes one or more persistent IP addresses from a cluster.

Syntax

```
smcli sysmirror/rmpersistentip [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmpersistentip {-c|--cluster} <CLUSTER> \
    [-C|--confirm] \
    <PERSISTENT_IP>[,<PERSISTENT_IP#2>,...]
smcli sysmirror/rmpersistentip {-c|--cluster} <CLUSTER> \
    [-C|--confirm] \
    {-a|--all}
```

Command alias

`rmipi`

Description

If you do not use the `--confirm` flag, a warning message is displayed asking you to use the `--confirm` flag.

Before you run this command, you must create a snapshot of the affected cluster for backup purposes.

The removal of the persistent IP address occurs on the controlling node of the specified cluster. None of the other nodes in the cluster are aware the removal of the persistent IP address. To complete the removal of the persistent IP address, you must synchronize the cluster.

This synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

`-a|--all`

Indicates that this operation must be performed on all instances of this class.

`-c|--clusters`

Displays one or more clusters on which to run this command.

- C|--confirm**
Provides confirmation to perform the operation.
- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

- <PERSISTENT_IP>**
A label that you can use to identify the persistent IP address.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Remove a persistent IP address that has a label of `const1` from `clusterA`:
`smcli sysmirror/rmpi --confirm -c clusterA const1`
2. Remove a persistent IP address of `10.14.187.4` from `clusterA`:
`smcli sysmirror/rmpi --confirm -c clusterA 10.14.187.4`
3. Remove all persistent IP addresses that are defined within `clusterA`:

Note: You must run this command only after you have planned for the removal. If you remove the fallback timer, your environment is affected.

```
smcli sysmirror/rmpi -C -c clusterA --all
```

rmresgrp command

Purpose

Removes the specified resource groups from the node. The resources that are managed by the specified resource groups might be unavailable during the remove process.

Syntax

```
smcli sysmirror/rmresgrp [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmresgrp {-c|--cluster} <CLUSTER> \
[-C|--confirm] \
```

```

smcli sysmirror/rmresgrp <RG>[,<RG#2>,...]
                        {-c|--cluster} <CLUSTER> \
                        [-C|--confirm] \
                        {-a|--all}

```

Command alias

rmrg

Description

Completely removes one or more networks from the specified cluster. If confirmation of the removal is not given by using the `--confirm` flag, the removal does not occur. You should create a snapshot of the affected cluster before running this command.

This new definition initially exists only on the controlling node of the specified cluster. To complete the creation, you must synchronize the cluster. This is a separate step, so that you can make multiple changes to the cluster, and then verify those changes as a set before enabling them by synchronization. You can discard your changes at any time before the synchronization occurs, and before changing the controlling node of the cluster.

Flags

- a|--all**
Indicates that this operation should be performed on all instances of this class.
- c|--cluster**
A cluster on which this command runs.
- C|--confirm**
Provides confirmation that performing this operation is desired.
- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

- <RG>**
The label of a resource group on which the operation is performed.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.

- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. To remove a resource group named `rg1` from a cluster named `clusterA`, enter:
`smcli sysmirror/rmresgrp --confirm -c clusterA rg1`
2. To remove all resource groups that are defined within cluster `clusterA`, enter:
`smcli sysmirror/rmrg -C -c clusterA --all`

Note: You should only run this command after you have planned for the removal and are aware of how the removal will affect your environment.

rmserviceip command

Purpose

Removes one or more service IP addresses from a cluster.

Syntax

```
smcli sysmirror/rmserviceip [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmserviceip {-c|--cluster} <CLUSTER> \
                             [-C|--confirm] \
                             <SERVICE_IP>[,<SERVICE_IP#2>,...]
smcli sysmirror/rmserviceip {-c|--cluster} <CLUSTER> \
                             [-C|--confirm] \
                             {-a|--all}
```

Command alias

`rmsi`

Description

If confirmation of the removal is not displayed when you use the `--confirm` flag, the service IP address is not removed. Before you run this command, you should create a snapshot of the affected cluster for backup purposes.

The removal of the service IP address occurs on the controlling node of the specified cluster. None of the other nodes in the cluster are aware the removal of the service IP address. To complete the removal of the service IP address, you must synchronize the cluster.

This synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

- a|--all**
Indicates that this operation must be performed on all instances of this class.
- c|--clusters**
Displays one or more clusters on which to run this command.
- C|--confirm**
Provides confirmation to perform the operation.
- h|-?|--help**
Requests help for this command.

-v|--verbose
Requests maximum details in the displayed information.

Parameters

<SERVICE_IP>
A label that you can use to identify the service IP address.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Remove a service IP address that has a label of `svc1` from `clusterA`:
`smcli sysmirror/rmsi --confirm -c clusterA svc1`
2. Remove a service IP address of `10.14.187.4` from `clusterA`:
`smcli sysmirror/rmsi --confirm -c clusterA 10.14.187.4`
3. Remove all service IP addresses that are defined within `clusterA`:

Note: You must run this command only after you have planned for the removal. If you remove all the service IP addresses, your environment is affected.

```
smcli sysmirror/rmsi -C -c clusterA --all
```

rmsite command

Purpose

Removes one or more PowerHA SystemMirror sites from the specified cluster.

Syntax

```
smcli sysmirror/rmsite [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmsite {-c|--cluster} <CLUSTER> \
    [-C|--confirm] \
    <SITE>[,<SITE#2>,...]
smcli sysmirror/rmsite --cluster <CLUSTER> \
    [-C|--confirm] \
    {-a|--all}
```


Description

If you do not use the **--confirm** flag, a warning message is displayed asking you to use the **--confirm** flag.

This new definition initially exists only on the controlling node of the specified cluster. None of the other nodes in the cluster are aware of it yet. To complete the creation process, the cluster must be synchronized by using the **synccluster** command.

This synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

-a|--all

Indicates that this operation must be performed on all instances of this class.

-c|--clusters

Displays one or more clusters on which to run this command.

-C|--confirm

Provides confirmation to perform the operation.

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<SITE>

The label of a site on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Remove a site named mySite from a cluster named clusterA:
`smcli sysmirror/rmsite --confirm -c clusterA mySite`

- Remove all sites that are defined within cluster clusterA:

Note: You must run this command only after you have planned for the removal of the site.

```
smcli sysmirror/rmsite -c -c clusterA --all
```

rmsnapshot command

Purpose

Removes one or more snapshots from a cluster.

Syntax

```
smcli sysmirror/rmsnapshot [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmsnapshot {-c|--cluster} <CLUSTER> \
                             [-C|--confirm] \
                             <SNAPSHOT>[,<SNAPSHOT#2>,...]
smcli sysmirror/rmsnapshot {-c|--cluster} <CLUSTER> \
                             [-C|--confirm] \
```

Command alias

rmss

Description

If you do not use the **--confirm** flag, a warning message is displayed asking you to use the **--confirm** flag.

Before you run this command, you must create a snapshot of the affected cluster for backup purposes.

Flags

- a|--all**
Indicates that this operation must be performed on all instances of this class.
- c|--clusters**
Displays one or more clusters on which to run this command.
- C|--confirm**
Provides confirmation to perform the operation.
- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

<SNAPSHOT>
A label that you can use to identify a snapshot.

Exit status

- 0** Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.

- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Remove a snapshot with a label `snap1` from the controlling node of `clusterA`:
`smcli sysmirror/rmsnapshot --confirm -c clusterA snap1`
2. Remove all snapshots that are defined on the controlling node of `clusterA`:
`smcli sysmirror/rmss -C -c clusterA --all`

rmtape command

Purpose

Removes one or more tape device definitions from a cluster.

Syntax

```
smcli sysmirror/rmtape [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmtape {-c|--cluster} <CLUSTER> \
                        [-C|--confirm] \
                        <TAPE>[,<TAPE#2>,...]
smcli sysmirror/rmtape {-c|--cluster} <CLUSTER> \
                        [-C|--confirm] \
                        {-a|--all}
```

Command alias

`rmtp`

Description

If you do not use the `--confirm` flag, a warning message is displayed asking you to use the `--confirm` flag.

Before you run this command, you should create a snapshot of the affected cluster for backup purposes.

The removal of the tape device definition occurs on the controlling node of the specified cluster. None of the other nodes in the cluster are aware the removal of the tape device definition. To complete the removal of the tape device definition, you must synchronize the cluster

This synchronization is done as a separate step so that you can make multiple changes to the cluster and then verify those changes. If you do not change the controlling node in the cluster, you can discard your changes at any time before the synchronization occurs.

Flags

`-a|--all`

Indicates that this operation must be performed on all instances of this class.

- c|--clusters**
Displays one or more clusters on which to run this command.
- C|--confirm**
Provides confirmation to perform the operation.
- h|-?|--help**
Requests help for this command.
- v|--verbose**
Requests maximum details in the displayed information.

Parameters

<TAPE>

The label that you can use to identify a tape device.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Remove the tape device definition `tape1` from the controlling node of `clusterA`:

```
smcli sysmirror/rmtape --confirm -c clusterA tape1
```
2. Remove all tape device definitions from the controlling node of `clusterA`:

```
smcli sysmirror/rmtape -C -c clusterA --all
```

rmuser command

Purpose

Removes one or more AIX users from the specified cluster.

Syntax

```
smcli sysmirror/rmuser [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rmuser {-c|--cluster} <CLUSTER> \
  {-R|--registry} <ldap|files>] \
  [-l|--leave_auth_info] \
  [{-r|--resgrp} <RESOURCE_GROUP>] \
  [-C|--confirm] \
  <user_name>[,<user_name#2>,...]
```

Command alias

rmur

Description

If you do not use the **--confirm** flag, a warning message is displayed asking you to use the **--confirm** flag.

Before you run this command, you must create a snapshot of the affected cluster for backup purposes.

Flags

-c|--clusters

Displays one or more clusters on which to run this command.

-C|--confirm

Provides confirmation to perform the operation.

-h|-?|--help

Requests help for this command.

-l|--leave_auth_info

Indicates that when the user is removed, any existing authorization information is not removed. By default behavior all authorization data is removed when the user is deleted.

-R|--registry

You can use one the following options to indicate from where to remove the user group definition from:

ldap Removes all information and credentials that are stored in an LDAP server.

files Removes all information and credentials that are stored locally in the cluster.

-r|--resgrp

Specify a resource group to ensure that the specified groups are removed from all the nodes that are defined within the resource group.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<user_name>

A label that you can use to identify an AIX user.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

1 The command has not been used in the intended manner.

2 A specified or required file was not found.

11 The command failed to accomplish its task.

12 A specified or a required resource was not found.

13 The required command input was not provided.

14 The command input was not valid.

- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. From the nodes that manage resource group RG1, remove a local user named hsimpson from clusterA:
`smcli sysmirror/rmuser --confirm -R files -r RG1 -c clusterA hsimpson`
2. Remove an LDAP group named msimpson from clusterA:
`smcli sysmirror/rmuser --confirm -R ldap -c clusterA msimpson`

rssnapshot command

Purpose

Restores the previously saved cluster configuration information from a backup file.

Syntax

```
smcli sysmirror/rssnapshot [-h|-?|--help] [-v|--verbose]
smcli sysmirror/rssnapshot {-c|--cluster} <CLUSTER> \
    [-C|--confirm] \
    [-u|--unconfigure] \
    [-F|--force] \
    <SNAPSHOT>
```

Command alias

rsss

Description

The configuration information that is saved in a snapshot is verified to check for any errors. If no errors are found, the snapshot is restored to the controlling node. After the snapshot is restored, the restored cluster data is synchronized from the controlling node to all the other nodes that were defined in the snapshot. The snapshot restoration process completely overwrites the existing cluster with the cluster data that the snapshot contains. However, a new snapshot of the existing cluster is automatically created before the restoration process begins. Thus, the old cluster can be restored if needed, thereby avoiding the possibility of losing cluster configuration data.

If an error is detected at any point during the restoration process, the process is stopped and the original existing cluster is not removed or changed.

Flags

-c|--clusters

Displays one or more clusters on which to run this command.

-C|--confirm

Provides confirmation to perform the operation.

-F|--force

Forces the snapshot to be restored, even if a verification error is detected. Use this flag if you know why the verification errors are occurring.

-h|-?|--help

Requests help for this command.

-u|--unconfigure

Changes a clusters configuration without affecting a running cluster. Changes are applied to the cluster the next time you verify and synchronize the cluster.

Note: Whenever a change is made to the cluster configuration, those changes are first verified, and then propagated to all nodes in the cluster during verification and synchronization. If changes are made when cluster services are active, those changes are reflected in the running cluster. For example, if you remove a resource group which manages an application, the resources for that application are taken offline when the changes are synchronized.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<SNAPSHOT>

A label that you can use to identify a snapshot.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Restore snapshot snap1 on the controlling node in clusterA:
smcli sysmirror/rssnapshot --confirm -c clusterA snap1
2. Restore snapshot snap2 on the controlling node in clusterA, which forces the restoration regardless of any errors:
smcli sysmirror/rsss --confirm -c clusterA -F snap2

startcluster command

Purpose

Brings all the nodes in a cluster to an active state.

Syntax

```
smcli sysmirror/startcluster [-h|-?|--help] [-v|--verbose]
smcli sysmirror/startcluster [{-w|--when} {now|restart|both}] \
    [-m|--manage] \
    [-b|--broadcast] \
    [-C|--start_clinfo] \
    [-g|--start_clinfo_cg] \
```

```
[-F|--force] \
[-x|--fix_errors] \
[{-t|--timeout} ##] \
<CLUSTER>[,<CLUSTER#2>,...]
```

Command alias

stcl

Description

Starts cluster services on all nodes in the specified cluster. All nodes that are moved to an active state are available to keep resource highly available.

Flags

-b|--broadcast

Indicates if a terminal announcement is sent to inform users of the start of cluster services. All users logged into the cluster will see this message.

-C|--start_clinfo

Indicates that the clinfoES subsystem is started when cluster services are started.

-g|--start_clinfo_cg

Indicates that the clinfoES subsystem is started with the consistency group support.

-F|--force

Attempts to force the cluster to start. Verification error, if any, are ignored.

-h|-?|--help

Requests help for this command.

-m|--manage

Informs PowerHA SystemMirror how to monitor and manage applications. This flag brings resource groups online according to their configuration settings and the current cluster state. PowerHA SystemMirror monitors the resource groups and application availability.

Note: If you do not use this flag, no action is taken on any resource groups.

-t|--timeout

The number of seconds to wait for the operation to complete. Use this flag to specify when to collect results even if the timeout occurs and an error is reported. The operation might still complete successfully.

-v|--verbose

Requests maximum details in the displayed information.

-w|--when

Select one of the following options to specify when to start PowerHA SystemMirror services:

now Start PowerHA SystemMirror services immediately.

restart Start PowerHA SystemMirror services when the AIX operating system starts up.

both Start PowerHA SystemMirror services by using the **restart** and **now** options.

-x|--fix_errors

PowerHA SystemMirror attempts to correct any reported verification errors. However, not all errors can be corrected automatically.

Parameters

<CLUSTER>

Specifies the label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Start clusterA:
`smcli sysmirror/startcluster clusterA`
2. Start clusterA immediately, and broadcast that start of the cluster. Also, automatically manage all resource groups and ally correct any verification problems that might be found:
`smcli sysmirror/stcl -x -b -m -w now clusterA`

startnode command

Purpose

Brings one or more nodes in a cluster to an active state.

Syntax

```
smcli sysmirror/startnode [-h|-?|--help] [-v|--verbose]
smcli sysmirror/startnode {-c|--cluster} <CLUSTER> \
    [{-w|--when} {now|restart|both}] \
    [-m|--manage] \
    [-b|--broadcast] \
    [-C|--start_clinfo] \
    [-g|--start_clinfo_cg] \
    [-F|--force] \
    [-x|--fix_errors] \
    [{-t|--timeout} ##] \
    <NODE>[,<NODE#2>,...]
```

Command alias

stnd

Description

Brings a node to an active state, by making it ready to keep resources highly available. Depending on the settings and previously selected policies, PowerHA SystemMirror might attempt to activate or acquire various resources that are associated with the node while it is being brought online.

Flags

-b|--broadcast

Indicates if a terminal announcement is sent to inform users of the start of cluster services. All users logged into the cluster will see this message.

-c|--clusters

Displays one or more clusters on which to run this command.

-C|--start_clinfo

Indicates that the clinfoES subsystem is started when cluster services are started.

-g|--start_clinfo_cg

Indicates that the clinfoES subsystem is started with the consistency group support.

-F|--force

Attempts to force the cluster to start. Verification error, if any, are ignored.

-h|-?|--help

Requests help for this command.

-m|--manage

Informs PowerHA SystemMirror how to monitor and manage applications. This flag brings resource groups online according to their configuration settings and the current cluster state. PowerHA SystemMirror monitors the resource groups and application availability.

Note: If you do not use this flag, no action is taken on any resource groups.

-t|--timeout

The number of seconds to wait for the operation to complete. Use this flag to specify when to collect results even if the timeout occurs and an error is reported. The operation might still complete successfully.

-v|--verbose

Requests maximum details in the displayed information.

-w|--when

Select one of the following options to specify when to start PowerHA SystemMirror services:

now Start PowerHA SystemMirror services immediately.

restart Start PowerHA SystemMirror services when the AIX operating system starts up.

both Start PowerHA SystemMirror services by using the **restart** and **now** options.

-x|--fix_errors

PowerHA SystemMirror attempts to correct any reported verification errors. However, not all errors can be corrected automatically.

Parameters

<NODE>

A label that you can use to identify a node.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

1 The command has not been used in the intended manner.

2 A specified or required file was not found.

- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Start nodeA:

```
smcli sysmirror/startnode nodeA
```
2. Start nodeA and nodeB immediately and broadcast a message about their launch. Also, automatically manage all resource groups and if possible, automatically correct any verification problems:

```
smcli sysmirror/stnd -x -b -m -w now nodeA,nodeB
```

startresgrp command

Purpose

Brings one or more PowerHA SystemMirror resource groups online.

Syntax

```
smcli sysmirror/startresgrp [-h|-?|--help] [-v|--verbose]
smcli sysmirror/startresgrp {-c|--cluster}<CLUSTER> \
    [{-n|--nodes} <NODE>[,<NODE#2>,...] ] \
    <RG>[,<RG#2>,...]
```

Command alias

strg

Description

Moves a resource group to the ONLINE state. PowerHA SystemMirror activates all the resources associated with the resource group while attempting to bring the resource group online.

Flags

-c|--cluster

The cluster on which to run this command.

-h|-?|--help

Requests help for this command.

-n|--node

One or more nodes on which to bring the resource group online. This option is typically only useful with concurrent resource groups.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<RG>

The label of a resource group on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1 The command has not been used in the intended manner.
- 2 A specified or required file was not found.
- 11 The command failed to accomplish its task.
- 12 A specified or a required resource was not found.
- 13 The required command input was not provided.
- 14 The command input was not valid.
- 15 A variable that this command depends on to run was not found.
- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. Start a resource group named db2grp on node nodeD:

```
smcli sysmirror/startresgrp -c clusterA \  
                             -n nodeD \  
                             db2grp
```
2. Start a concurrent resource group named oracle_grp on nodes nodeA, nodeB, and nodeC:

```
smcli sysmirror/startresgrp --cluster clusterA \  
                             --node nodeA,nodeB,nodeC \  
                             oracleGrp
```

stopcluster command

Purpose

Brings all the nodes in a clusters to an inactive state.

Syntax

```
smcli sysmirror/stopcluster [-h|-?|--help] [-v|--verbose]  
smcli sysmirror/stopcluster [{-w|--when} {now|restart|both}] \  
                             [-u|--unmanage] \  
                             [-b|--broadcast] \  
                             [{-t|--timeout} ##] \  
                             <CLUSTER>[,<CLUSTER#2>,...]
```

Command alias

spcl

Description

Stops cluster services on all nodes in the specified cluster. When nodes are in an inactive state, all resources that are managed by this cluster might not be highly available.

Flags

-b|--broadcast

Indicates if a terminal announcement is sent to inform users of the start of cluster services. All users logged into the cluster will see this message.

-h|-?|--help

Requests help for this command.

-u|--unmanage

Indicates that after the nodes is inactive, the resources managed by the cluster remain active.

-t|--timeout

The number of seconds to wait for the operation to complete. Use this flag to specify when to collect results even if the timeout occurs and an error is reported. The operation might still complete successfully.

-v|--verbose

Requests maximum details in the displayed information.

-w|--when

Select one of the following options to specify when to start PowerHA SystemMirror services:

now Start PowerHA SystemMirror services immediately.

restart Start PowerHA SystemMirror services when the AIX operating system starts up.

both Start PowerHA SystemMirror services by using the **restart** and **now** options.

Parameters

<CLUSTER>

Specifies the label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Stop clusterA:
smcli sysmirror/stopcluster clusterA
2. Stop clusterA immediately and leave all resources that were being managed by the resource groups in an active state:
smcli sysmirror/stc1 --when now --unmanage clusterA

stopnode command

Purpose

Brings one or more nodes in a cluster to an inactive state.

Syntax

```
smcli sysmirror/stopnode [-h|-?|--help] [-v|--verbose]
smcli sysmirror/stopnode {-c|--cluster} <CLUSTER> \
    [{-w|--when} {now|restart|both}] \
    [{-m|--manage} {offline|move|unmanage}] \
    [-b|--broadcast] \
    [{-t|--timeout} ##] \
    <NODE>[,<NODE#2>,...]
```

Command alias

spnd

Description

Stops cluster services on all nodes in the specified cluster. When nodes are in an inactive state, all resources that are managed by this cluster might not be highly available.

Depending on the chosen settings and previously selected policies, PowerHA SystemMirror might attempt to either deactivate or move various resources that are associated with the node while it is being brought offline.

Flags

-b|--broadcast

Indicates if a terminal announcement is sent to inform users of the start of cluster services. All users logged into the cluster will see this message.

-c|--clusters

Displays one or more clusters on which to run this command.

-h|-?|--help

Requests help for this command.

-m|--manage

Informs PowerHA SystemMirror how to monitor and manage applications. This flag brings resource groups online according to their configuration settings and the current cluster state. PowerHA SystemMirror monitors the resource groups and application availability.

Note: If you do not use this flag, no action is taken on any resource groups.

-t|--timeout

The number of seconds to wait for the operation to complete. Use this flag to specify when to collect results even if the timeout occurs and an error is reported. The operation might still complete successfully.

-v|--verbose

Requests maximum details in the displayed information.

-w|--when

Select one of the following options to specify when to start PowerHA SystemMirror services:

now Start PowerHA SystemMirror services immediately.

restart Start PowerHA SystemMirror services when the AIX operating system starts up.

both Start PowerHA SystemMirror services by using the **restart** and **now** options.

Parameters

<NODE>

A label that you can use to identify a node.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Stop nodeA:

```
smcli sysmirror/stopnode nodeA
```
2. Stop nodeA and nodeB immediately and broadcasting their shutdown. Also, stop managing any cluster-managed resources that are currently active on those nodes:

```
smcli sysmirror/spnd -b -m unmanage -w now nodeA,nodeB
```

stopresgrp command

Purpose

Brings one or more PowerHA SystemMirror resource groups offline.

Syntax

```
smcli sysmirror/stopresgrp [-h|-?|--help] [-v|--verbose]
smcli sysmirror/stopresgrp {-c|--cluster}<CLUSTER> \
    [{-n|--nodes} <NODE>[,<NODE#2>,...] ] \
    <RG>[,<RG#2>,...]
```

Command alias

sprg

Description

Moves a resource group to the OFFLINE state. PowerHA SystemMirror stops all the resources that are associated with the resource group while attempting to bring the resource group offline.

Flags

-c|--cluster

The cluster on which to run this command.

-h|-?|--help

Requests help for this command.

-n|--nodes

One or more nodes on which to bring the resource group offline. This option is typically only useful with concurrent resource groups.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<RG>

The label of a resource group on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Stop a resource group named db2grp on node nodeD:

```
smcli sysmirror/stopresgrp -c clusterA \  
                           -n nodeD \  
                           db2grp
```

2. Stop a concurrent resource group named oracle_grp on nodes nodeA, nodeB, and nodeC:

```
smcli sysmirror/stopresgrp --cluster clusterA \  
                           --node nodeA,nodeB,nodeC \  
                           oracleGrp
```

synccluster command

Purpose

Verifies and synchronizes a PowerHA SystemMirror cluster.

Syntax

```
smcli sysmirror/synccluster {-h|-?|--help} [-v|--verbose]
smcli sysmirror/synccluster {-n|--no_verification} \
<CLUSTER>
smcli sysmirror/synccluster [-f|--fix_errors] \
[-C|--changes_only] \
[-t|--custom_tests_only] \
[{-M|--methods} <METHOD>[,<METHOD#2>,...] ] \
[{-e|--maximum_errors} <##>] \
[-F|--force] \
[-V|--verbose_logging] \
[{-l|--logfile} <full_path_to_file>] \
<CLUSTER>
```

Command alias

sync

Description

Copies the cluster configuration from the controlling node of the specified cluster to each of the other nodes in the cluster. If cluster services are running on any node, the configuration changes take effect, possibly causing one or more resources to change state.

Flags

-h|-?|--help

Requests help for this command.

-C|--changes_only

This flag will cause only those verification tests to be run that directly relate to the PowerHA SystemMirror configuration information that you have changed. By default, all verification tests are run.

-e|--maximum_errors

By default, the program continues to run after an error is encountered, allowing a full list of errors to be generated. If an error count is specified, the program ends after encountering the specified number of errors.

-F|--force_sync

Forces the synchronization to occur even if verification errors are detected. This operations must only be done with great care and expert knowledge of why the verification errors are occurring.

-f|--fix_errors

This option causes PowerHA SystemMirror to attempt to correct any reported verification errors. However, not all errors can be corrected automatically.

-l|--logfile

Specifies a log file in which to store the verification output into.

-M|--methods

The labels of one or more custom verification methods to call as part of the verification process. If the special value **All** is specified all currently defined verification methods are run.

-n|--no_verification

Indicates that no verification is performed. In general, do not specify this flag prior to performing a synchronization.

-t|--custom_tests_only

This flag causes only the specified custom verification methods to be run. None of the built-in tests are called.

-v|--verbose_logging

Requests maximum details in the displayed information, displaying all the information that normally is written to the **clverify.log** file.

Parameters

<CLUSTER>

The label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Synchronize a cluster named `clusterA` in the simplest manner possible:

```
smcli sysmirror/synccluster clusterA
```
2. Synchronize a cluster named `clusterB`, specifying two custom verification methods named `checkAppA` and `checkDiskSpace`:

```
smcli sysmirror/synccluster -m checkAppA,checkDiskSpace \  
clusterB
```
3. Synchronize a cluster named `clusterC`, indicating that PowerHA SystemMirror should attempt to correct any verification errors that it finds, but abort the command if any errors cannot be corrected:

```
smcli sysmirror/synccluster -m checkAppA,checkDiskSpace \  
--fix_errors \  
--maximum_errors 1 \  
clusterC
```

syncfilecollection command

Purpose

Copies files and directories from the controlling node to the other nodes in the cluster.

Syntax

```
smcli sysmirror/ syncfilecollection [-h|-?|--help] [-v|--verbose]  
smcli sysmirror/ syncfilecollection {-c|--cluster} <CLUSTER> \  
<FC>
```

Command alias

`syfc`

Description

Files and directories in the specified file collection are copied from the controlling node to all the other nodes in the cluster. Any existing files and directories on the other nodes are overwritten.

Flags

-c|--clusters

Displays one or more clusters on which to run this command.

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<FC>

A label that you can use to identify a file collection.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Example

The following example synchronizes a file collection named fc1 on clusterA:

```
smcli sysmirror/syfc -c clusterA fc1
```

undochanges command

Purpose

Restores the last synchronized cluster configuration on the controlling node.

Syntax

```
smcli sysmirror/undochanges {-h|-?|--help} [-v|--verbose]  
smcli sysmirror/undochanges <CLUSTER>
```

Command alias

undo

Description

This operation restores the PowerHA SystemMirror Configuration Database, from the Active Configuration Database. Typically, this has the effect of discarding any unsynchronized changes. Before proceeding with this operation, a snapshot named UserModifiedDB is created automatically, so you regain the changes you just discarded.

Flags

-h|-?|--help
Requests help for this command.

-v|--verbose
Requests maximum details in the displayed information.

Parameters

<CLUSTER>
The label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Discard any unsynchronized changes on a cluster named clusterA:
`smcli sysmirror/undochanges clusterA`

vercluster command

Purpose

Verifies the configuration of an PowerHA SystemMirror cluster.

Syntax

```
smcli sysmirror/vercluster {-h|-?|--help} [-v|--verbose]
smcli sysmirror/vercluster [-f|--fix_errors] \
                             [-C|--changes_only] \
```

```

[-t|--custom_tests_only] \
[{-M|--methods} <METHOD>[,<METHOD#2>,...] ] \
[{-e|--maximum_errors} <##>] \
[-F|--force] \
[-V|--verbose_logging] \
[{-l|--logfile} <full_path_to_file>] \
<CLUSTER>
smcli sysmirror/vercluster {-s|--synchronize} \
[-F|--force] \
[-f|--fix_errors] \
[-C|--changes_only] \
[-t|--custom_tests_only] \
[{-M|--methods} <METHOD>[,<METHOD#2>,...] ] \
[{-e|--maximum_errors} <##>] \
[-F|--force] \
[-V|--verbose_logging] \
[{-l|--logfile} <full_path_to_file>] \
<CLUSTER>

```

Command alias

vrcl

Description

Verifies the validity of the cluster configuration on the controlling node of the specified cluster.

Flags

-h|-?|--help

Requests help for this command.

-C|--changes_only

This flag causes only those verification tests to be run that directly relate to the PowerHA SystemMirror configuration information that you changed. By default, all verification tests are run.

-e|--maximum_errors

By default, the program continues to run after an error is encountered, allowing a full list of errors to be generated. If an error count is specified, the program terminates after encountering the specified number of errors.

-F|--force_sync

Forces the synchronization to occur even if verification errors are detected. This operation must be done with great care and expert knowledge of why the verification errors are occurring.

-f|--fix_errors

This option causes PowerHA SystemMirror to attempt to correct any reported verification errors. However, not all errors can be corrected automatically.

-l|--logfile

Specifies a log file in which to store the verification output into.

-M|--methods

The labels of one or more custom verification methods to call as part of the verification process. If the special value of **All** is specified, all currently defined verification methods are run.

-s|--synchronize

Causes the cluster configuration from the controlling node to be copied to each of the other nodes in the cluster.

-t|--custom_tests_only

This flag causes only the specified custom verification methods to be run. None of the built-in tests are invoked.

-v|--verbose_logging

Requests maximum details in the displayed information, displaying all the information that normally is written to the **clverify.log** file.

Parameters

<CLUSTER>

The label of a cluster on which to perform this operation.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. Verify a cluster named `clusterA` in the simplest manner possible:
`smcli sysmirror/vercluster clusterA`
2. Verify and synchronize a cluster named `clusterB`:
`smcli sysmirror/vercluster -s clusterB`
3. Verify and synchronize a cluster named `clusterC`, using two custom verification methods, `checkAppA` and `checkDiskSpace`, indicating that PowerHA SystemMirror should attempt to correct any verification errors that it finds, but aborting the command if any errors cannot be corrected:

```
smcli sysmirror/vercluster -M checkAppA,checkDiskSpace \  
    --synchronize \  
    --fix_errors \  
    --maximum_errors 1 \  
    clusterC
```

vlog command

Purpose

Displays the contents of a log file.

Syntax

```
smcli sysmirror/vlog [-h|-?|--help] [-v|--verbose]  
smcli sysmirror/vlog {-c|--cluster} <CLUSTER> \  
    [{-n|--node} <NODE>] \  
    [{-T|--tail} ##] \  
    [{-F|--filters} <PATTERN>[,<PATTERN#2>,...]] \  
    <LOG>
```

Description

You can use log files to troubleshoot PowerHA SystemMirror problems. Use this command to access the logs files that are located on any node in the cluster. The default settings for this command displays the log files for the controlling node.

Flags

-c|--clusters

Displays one or more clusters on which to run this command.

-F|--filters

One or more search expressions that are separated by a delimiter. The default delimiter is a comma.

-h|-?|--help

Requests help for this command.

-n|--node

Identifies the node that displays the log files.

-T|--tail

The number of lines to display from the end of the output.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<LOG>

A label that you can use to identify the log.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.
- 16** A specified search failed to find any matching data.
- 29** The current locale is not supported.

Examples

1. View a log file named `hacmp.out` on the controlling node in `clusterA`:
`smcli sysmirror/vlog -c clusterA hacmp.out`
2. View a log file named `clutils.log` on `nodeC` in `clusterA` that only returns lines from the log that contain `ERROR`:
`smcli sysmirror/vlog -n nodeC -c clusterA -F "ERROR" clutils.log`

vreport command

Purpose

Provides a variety of reports about one or more clusters.

Syntax

```
smcli sysmirror/vreport [-h|-?|--help] [-v|--verbose]
smcli sysmirror/vreport {-c|--clusters} <CLUSTER>[,<CLUSTER#2>,...] \
    <REPORT>
smcli sysmirror/vreport {-a|--all} \
    <REPORT>
```

Command alias

vwrp

Description

Several reports are available for PowerHA SystemMirror from its command line interface on the IBM Systems Director server. These reports can be viewed and/or stored for historical records.

Flags

-a|--all

Indicates that this operation should be performed on all instances of this class.

-c|--clusters

One or more clusters to run this command on.

-h|-?|--help

Requests help for this command.

-v|--verbose

Requests maximum details in the displayed information.

Parameters

<REPORT>

A report to be generated and displayed. The report must be one of the following types: applications, availability, basic, cluster, dependencies, events, fsinfo, lvinfo, networks, nodeinfo, rginfo, status, topology, or vginfo.

Exit status

0 Indicated that the command completed successfully.

A nonzero value indicates a problem condition. Several of those conditions have exit values to assist you in problem determination.

- 1** The command has not been used in the intended manner.
- 2** A specified or required file was not found.
- 11** The command failed to accomplish its task.
- 12** A specified or a required resource was not found.
- 13** The required command input was not provided.
- 14** The command input was not valid.
- 15** A variable that this command depends on to run was not found.

- 16 A specified search failed to find any matching data.
- 29 The current locale is not supported.

Examples

1. View the basic report on cluster clusterA:
`smcli sysmirror/vreport -c clusterA basic`
2. View the status report on all PowerHA SystemMirror clusters that are currently managed by IBM Systems Director:
`smcli sysmirror/vreport --all status`

Notices

This information was developed for products and services that are offered in the USA.

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

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

*IBM Corporation
Dept. LRAS/Bldg. 903
11501 Burnet Road
Austin, TX 78758-3400
USA*

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

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

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

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

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

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

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

COPYRIGHT LICENSE:

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

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

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

© Copyright IBM Corp. _enter the year or years_. All rights reserved.

Privacy policy considerations

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

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

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

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

Trademarks

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

INFINIBAND, InfiniBand Trade Association, and the INFINIBAND design marks are trademarks and/or service marks of the INFINIBAND Trade Association.

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

Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

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

Index

A

Adding

- application controller 28
- application monitor 29
- Capacity on Demand 30
- disk to an LVM mirror pool 34
- file collections 22
- interface 12
- network 12
- nodes to a site 42
- repository disk 39
- security user 16

Alerts

- creating 20
- testing 20
- viewing 20
- viewing event subscriptions 21

Application controller

- adding 28
- changing properties 28
- creating 27

Application monitor

- adding 29
- changing properties 29
- creating 28

C

Capacity on Demand

- adding 30
- changing properties 30

Changing 19, 36

- application controller properties 28
- application monitor properties 29
- Capacity on Demand properties 30
- custom verification method 46
- encrypted file system 18
- event method properties 43
- name of a site 41
- security group properties 15
- security user password 17
- security user properties 17
- storage agent 48
- storage system 48
- user's password 19

cluster properties 36

Cluster simulator 9

- using offline mode 9
- using planning mode 10

Configuring

- report settings 24

Creating 19

- alerts 20
- application controller 27
- application monitor 28
- cluster 11
- custom snapshot method 14
- custom verification method 46
- event command 44
- event method report 44

Creating (continued)

- file collections 21
- file system 33
- LDAP connection 17
- LVM mirror pool 33
- reports 24
- resource group dependencies 34
- security group 15
- security user 16
- snapshot 13
- storage agent 47
- storage system 48

Custom verification method

- changing 46
- creating 46
- viewing 45

D

Deleting

- event method 45

E

Encrypted file system

- changing 18
- setting up 18

Event command

- creating 44
- viewing 44

Event method

- changing properties 43
- creating report 44
- deleting 45
- viewing 42

F

File collections

- adding 22
- creating 21
- synchronizing 22
- viewing 21

I

Installing 4

- on nodes 6
- on the server 5
- prerequisites 5
- unlocking nodes 7

L

LDAP connection

- creating 17
- disconnecting 18

Log

- collecting log data 24

Log *(continued)*
 viewing data 23
 viewing files 23
Logical Volume Manager
 adding disk 34
 creating a mirror pool 33

R

RBAC roles 19
 changing 19
 creating 19
Report
 configuring report settings 24
 creating 24
 emailing 25
 uploading 25
Repository disk
 adding 39
 viewing 40
Resource group dependencies
 creating 34
 viewing 34
Restoring
 snapshot 14, 35

S

SAP liveCache Hot Standby 11
Security group
 adding a security user 16
 changing properties 15
 creating 15
Site
 adding nodes 42
 changing name 41
 site services 42
Snapshot
 creating 13
 creating custom method 14
 restoring 14
 viewing 14
Split or merge
 manual response 38
Storage agent
 changing 48
 creating 47
Storage system
 changing 48
 creating 48
Synchronizing
 file collections 22

U

Using
 offline simulation mode 9
 planning mode 10

V

Viewing
 alerts 20
 alerts event subscriptions 21
 custom verification method 45

Viewing *(continued)*
 event command 44
 event method 42
 file collections 21
 log data 23
 log files 23
 repository disk 40
 resource group dependencies 34
 resources 26
 snapshot 14



Printed in USA