


IBM PowerHA SystemMirror for AIX

Standard Edition

Version 7.2.1

Installing PowerHA SystemMirror




IBM PowerHA SystemMirror for AIX

Standard Edition

Version 7.2.1

Installing PowerHA SystemMirror



Note

Before using this information and the product it supports, read the information in "Notices" on page 77.

This edition applies to IBM PowerHA SystemMirror 7.2.1 Standard Edition for AIX and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2016, 2017.**

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

Contents

About this document	v
Highlighting	v
Case-sensitivity in AIX	v
ISO 9000.	v
Related information	v
Installing PowerHA SystemMirror	1
What's new in Installing PowerHA SystemMirror	1
Overview of the installation process	1
Creating shared LVM components	2
Logical volumes	3
Configuring PowerHA SystemMirror to use NFS version 4	4
Upgrading a PowerHA SystemMirror cluster	6
Migrating from PowerHA SystemMirror 6.1 to PowerHA SystemMirror 7.1.3 or PowerHA SystemMirror 7.2.0	6
Upgrading to PowerHA SystemMirror Version 7.2.1 prerequisites.	9
Planning for an upgrade	10
Upgrade do's and don'ts	12
Performing a rolling migration	14
Upgrading PowerHA SystemMirror using a snapshot	16
Upgrading an offline cluster for PowerHA SystemMirror.	17
Upgrade PowerHA SystemMirror using the cl_ezupdate command.	18
Additional migration tasks	19
Verifying the success of your migration	20
Troubleshooting your migration	22
Recovering your previous version of PowerHA SystemMirror.	22
Installing PowerHA SystemMirror on server nodes	23
Prerequisites for installing PowerHA SystemMirror on server nodes	23
Installation medium	25
PowerHA SystemMirror installation choices	26
Installing from an installation server	27
Installing from a hard disk	27
Installing from the installation medium	28
Completing the installation	30

Obtaining PowerHA SystemMirror APAR	30
Troubleshooting during the installation	30
Removing the PowerHA SystemMirror software	30
Installing PowerHA SystemMirror on client nodes	31
Step 1: Install the base system client images	31
Step 2: Copy the clhosts.client file onto client nodes	32
Step 3: Edit the clinfo.rc script	32
Step 4: Update the ARP cache for clients not using clinfo	33
Configuring installed hardware.	33
Configuring network Ethernet interface cards	33
Installing and configuring shared Fibre Channel tape drives	33
Configuring the installation of a shared tape drive	34
Defining shared LVM components.	35
Defining shared LVM components for nonconcurrent access	35
Defining LVM components for concurrent access	41
Configuring AIX for PowerHA SystemMirror	45
Networking considerations	45
Planning PowerHA SystemMirror file collections	46
Types of error notification	48
OEM disk, volume group, and file systems in a cluster	54
Integrating OEM disks in a PowerHA SystemMirror cluster	55
Integrating OEM volume groups in a PowerHA SystemMirror cluster	61
Integrating OEM file systems in a PowerHA SystemMirror cluster	66
PowerHA SystemMirror and SNMP utilities	71
PowerHA SystemMirror SNMP components	72
Systems monitor for AIX	74
Trap conflicts between SMUX peer daemons	74

Notices	77
Privacy policy considerations	79
Trademarks	79
Index	81

About this document

This document contains information about how to install and upgrade PowerHA[®] SystemMirror[®] for AIX[®].

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.

Case-sensitivity in AIX

Everything in the AIX operating system is case-sensitive, which means that it distinguishes between uppercase and lowercase letters. For example, you can use the **ls** command to list files. If you type **LS**, the system responds that the command is not found. Likewise, **FILEA**, **FiLea**, and **filea** are three distinct file names, even if they reside in the same directory. To avoid causing undesirable actions to be performed, always ensure that you use the correct case.

ISO 9000

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

Related information

- The PowerHA SystemMirror Version 7.2.1 PDF documents are available in the PowerHA SystemMirror 7.2.1 PDFs topic.
- The PowerHA SystemMirror Version 7.2.1 release notes are available in the PowerHA SystemMirror 7.2.1 release notes topic.

Installing PowerHA SystemMirror

You must understand concepts and planning information about PowerHA SystemMirror, before you install PowerHA SystemMirror.

What's new in Installing PowerHA SystemMirror

Read about new or significantly changed information for the Installing PowerHA SystemMirror topic collection.

How to see what's new or changed

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

June 2017

Added information about upgrading the PowerHA SystemMirror Version 7.2.1 by using the `cl_ezupdate` command in the new “Upgrade PowerHA SystemMirror using the `cl_ezupdate` command” on page 18 topic.

December 2016

Updated information about migrating to PowerHA SystemMirror Version 7.2.1 with existing split and merge policies in the “Upgrading to PowerHA SystemMirror Version 7.2.1 prerequisites” on page 9 topic.

Overview of the installation process

After completing the planning steps, you are ready to install the PowerHA SystemMirror software.

Step 1: Installing PowerHA SystemMirror on server nodes

In this step, you install PowerHA SystemMirror on all server nodes. Installing PowerHA SystemMirror on server nodes describes this step of the installation process.

Step 2: Installing PowerHA SystemMirror on client nodes

In this step, you install PowerHA SystemMirror on all client nodes. Installing PowerHA SystemMirror on client nodes describes this step of the installation process.

After installing PowerHA SystemMirror, you are ready to create a PowerHA SystemMirror cluster.

Overview of the cluster creation tasks

After installing the PowerHA SystemMirror software, you are ready to create a PowerHA SystemMirror cluster. This section provides an overview of the cluster creation process.

Step 1: Creating a PowerHA SystemMirror cluster

Note: Smart Assist products that provide rapid and flexible configuration for middleware solutions, such as SAP, DB2®, Websphere, Oracle, and other applications. For more information on the full range of available Smart Assist products, see Smart Assist for PowerHA SystemMirror.

Step 2: Configuring a PowerHA SystemMirror cluster

In this step, you perform tasks such as the following:

- Configuring installed hardware
- Defining shared Logical Volume Manager (LVM) components
- Configuring AIX for PowerHA SystemMirror

After you create a basic cluster and configure supporting components, you are ready to configure and monitor your cluster as described in the *Administering PowerHA SystemMirror*.

Overview of the migration tasks

If you are migrating an existing installation of PowerHA SystemMirror, follow the instructions from *Upgrading a PowerHA SystemMirror cluster*, and refer to *Installing PowerHA SystemMirror on server nodes* as needed.

Where to find documentation

The product README file and WhatsNewInThisRelease file are installed in the `/usr/es/sbin/cluster` directory.

Related concepts:

“Configuring installed hardware” on page 33

These topics describe how to ensure that network interface cards (NICs), shared external disk devices, and shared tape drives are ready to support a PowerHA SystemMirror cluster.

“Defining shared LVM components” on page 35

These topics describe how to define the LVM components shared by the nodes in a PowerHA SystemMirror cluster.

“Configuring AIX for PowerHA SystemMirror” on page 45

These topics discuss several general tasks necessary for ensuring that your PowerHA SystemMirror environment works as planned.

“Upgrading a PowerHA SystemMirror cluster” on page 6

These topics provide instructions for upgrading an existing PowerHA SystemMirror cluster configuration.

Related reference:

“Installing PowerHA SystemMirror on client nodes” on page 31

These topics describe how to install and configure the PowerHA SystemMirror software on client systems. They also describe how to edit files and scripts related to the Cluster Information (Cinfo) program.

Related information:

Planning PowerHA SystemMirror

Administering PowerHA SystemMirror

Creating shared LVM components

Setting up shared LVM components for a PowerHA SystemMirror cluster depends on the type of shared disk device and the method of shared disk access.

If you are planning to use OEM disks, volume groups, or file systems in your cluster (including Veritas volumes) see OEM disk, volume group, and file systems accommodation.

Prerequisites

At this point, you should have completed the planning steps described in the *Planning Guide*.

You should also be familiar with how to use the Logical Volume Manager (LVM). For information about AIX LVM, see the *AIX System Management guide*.

Related concepts:

“OEM disk, volume group, and file systems in a cluster” on page 54

These topics describe how you can customize PowerHA SystemMirror software to integrate original equipment manufacturer (OEM) disks, volume groups, and file systems in a PowerHA SystemMirror cluster.

Related information:

Planning PowerHA SystemMirror

Logical volumes

A *logical volume* is a set of logical partitions that AIX makes available as a single storage unit, that is, the logical view of a disk.

A *logical partition* is the logical view of a physical partition. Logical partitions can be mapped to one, two, or three physical partitions to implement mirroring.

In the PowerHA SystemMirror environment, logical volumes can be used to support a journaled file system or a raw device.

Specify the **superstrict** disk allocation policy for the logical volumes in volume groups for which forced varyon is specified. Using this configuration allows for the following properties:

- Guarantees that copies of a logical volume always reside on separate disks
- Increases the chances that forced varyon will be successful after a failure of one or more disks.

If you plan to use forced varyon for the logical volume, apply the **superstrict** disk allocation policy for disk enclosures in the cluster.

To specify the **superstrict** disk allocation policy, complete the following steps:

1. From the command line, enter `smit cl_admin`.
2. In SMIT, select **Storage > Logical Volumes** and select **Add a Logical Volume** or **Change a Logical Volume**.
3. Select **Allocate each logical partition copy on a separate physical volume**.
4. When using the **superstrict** disk allocation policy, specify the correct number of physical volumes for this logical volume. Do not use the default setting of 32 physical volumes.

Default NFS mount options for PowerHA SystemMirror

When mounting NFS file systems, PowerHA SystemMirror uses the default options `hard` and `intr`.

To set soft mounts or any other options when you are mounting an NFS file system, complete the following steps:

1. Enter `smit mknfsmnt`.
2. In the **MOUNT now, add entry to /etc/filesystems or both?** field, select the **file systems** option.
3. In the **/etc/filesystems entry will mount the directory on system RESTART** field, accept the default value of **no**.

This procedure adds the options you chose to the `/etc/filesystems` entry created. The PowerHA SystemMirror scripts then use any options you selected.

Creating and configuring NFS mount points on clients

An NFS mount point is required to mount a file system via NFS. In a nonconcurrent resource group, all the nodes in the resource group have been mounted using the NFS file system. The NFS mount point must be outside the directory tree of the local mount point.

After you create the NFS mount point on all nodes in the resource group, configure the **NFS File system to NFS Mount** attribute for the resource group.

To create NFS mount points and to configure the resource group for the NFS mount:

1. On each node in the resource group, create an NFS mount point by executing the following command: `mkdir /mountpoint`
where *mountpoint* is the name of the local NFS mount point over which the remote file system is mounted.
2. In the **Change/Show Resources and Attributes for a Resource Group** SMIT panel, the **File system to NFS Mount** field must specify both mount points.
Specify the NFS mount point, then the local mount point, separating the two with a semicolon. For example: `/nfspoint;/localpoint`
If there are more entries, separate them with a space: `/nfspoint1;/local1 /nfspoint2;/local2`
3. Optional: If there are nested mount points, nest the NFS mount points in the same manner as the local mount points so that they match up properly.
4. Optional: When cross-mounting NFS file systems, set the **File systems Mounted before IP Configured** field in SMIT for the resource group to **true**.

Configuring PowerHA SystemMirror to use NFS version 4

PowerHA SystemMirror supports NFS protocol Version 4 (NFS V4).

To ensure that PowerHA SystemMirror properly identifies NFS file systems mounted for NFS V4, complete the following steps:

1. Correctly set up NFS V4 configuration.
2. Make this configuration consistent on all nodes.

The fields needed to configure PowerHA SystemMirror to use NFS V4 are described in this section.

Step 1: Configuring NFS and changing to version 4

For PowerHA SystemMirror to recognize NFS V4, you change the NFS version on one node in the cluster in the AIX operating system first, and then on the rest of the nodes.

To change the NFS version on one node in the cluster:

1. Enter the fast path `smitty nfs`.
2. In SMIT, select **Network File System (NFS) > Configure NFS on This System > Change Version 4 Server Root Node** and press Enter.
3. Enter field values on the Change Version 4 Server Root Node panel as follows:
 - a. For the **Root Node Directory** field, enter the root node directory, for example `[/]`.
 - b. For the **Change number now, system restart, or both** field, select **both** from the picklist so that changes requested take place immediately and for every subsequent system restart.

You must also change the NFS version on each node in the cluster in the AIX operating system.

To change the NFS version on each node in the cluster:

1. Enter the fast path `smitty nfs`.

2. In SMIT, select **Network File System (NFS) > Configure NFS on This System > Change Version 4 Server Public Node** and press Enter.
3. Enter field values on the Change Version 4 Server Public Node panel as follows:
 - a. For the **Public Node Directory** field, enter the public node directory, for example [/].
 - b. For the **Change number now, system restart, or both** field, select **both** from the picklist so that changes requested take place immediately and for every subsequent system restart.

Step 2: Configuring the NFS local domain

Use SMIT to set the local domain on each node.

To set the local domain on each node using SMIT:

1. Enter the fast path `smitty nfs`
2. In SMIT, select **Network File System (NFS) > Configure NFS on This System > Configure NFS Local Domain > Change NFS Local Domain** and press Enter.
3. On the Display Current NFS Local Domain panel, for the **NFS Local Domain** field specify the new NFS local domain of the system.

Step 3: Exporting and mounting file systems

NFS exports are configured into PowerHA SystemMirror through resource groups. After NFS file systems are added to a resource group, PowerHA SystemMirror will actively manage and recover the file systems. During cluster events, PowerHA SystemMirror will export, unexport, cross-mount and mount the filesystem as needed, and will coordinate the recovery of the file system on a backup node.

The exports can be added and mounts can be specified in either of the following ways:

- Using NFS Configuration Assistant. This is designed to help you configure, view, change, or delete resource groups with NFS exports. The Configuration Assistant creates a new resource group with the specified NFS exports and mounts.
- Using the resource group **Change/Show Resources and Attributes for a Resource Group** panel. This is used to add, modify, or delete NFS exports and mounts to an already existing resource group.

Related information:

Using NFS with PowerHA SystemMirror

Configuring PowerHA SystemMirror cluster topology and resources (extended)

Step 4: Editing the `/usr/es/sbin/cluster/etc/exports` file

Modify the PowerHA SystemMirror `/usr/es/sbin/cluster/etc/exports` file on each node in the cluster to add the IP addresses for the network.

You can edit the file on one node and copy it to other cluster nodes. You can also use the PowerHA SystemMirror file collection function to keep this file in sync on all of the nodes of the cluster.

To modify the `/usr/es/sbin/cluster/etc/exports` file on each PowerHA SystemMirror cluster node, edit the `/usr/es/sbin/cluster/etc/exports` file on the control workstation using the following command:

```
vi /usr/es/sbin/cluster/etc/exports
```

For each file system, there should be a line that looks like this:

```
/fs/fs3big -vers=4,sec=sys:krb5p:krb5i:krb5:dh:none,rw,root=192.168.20.1:192.168.20.1:192.168.20.2:192.168.20.3:192.168.21.1:192.168.21.2:192.168.21.3:192.168.30.1:192.168.30.2:192.168.30.3
```

Note: To specify the Internet Protocol version 6 address, you can use the **-network** and **-umask** options in the `/usr/es/sbin/cluster/etc/exports` file.

Using this alternative exports file is optional. PowerHA SystemMirror checks the `/usr/es/sbin/cluster/etc/exports` file when Network File System (NFS) exports a file system or directory. If there is an entry for the file system or directory in this file, PowerHA SystemMirror uses the options listed. However, in some cases, PowerHA SystemMirror might ignore the version option as described in the Administration guide. If the file system or directory that you are trying to export with NSF is not listed in the file or if the alternate file does not exist, the file system or directory is exported with the default option of root access for all cluster nodes.

Related information:

Verifying and synchronizing a PowerHA SystemMirror cluster

Configuring PowerHA SystemMirror cluster topology and resources (extended)

Step 5: Modify entries in `/etc/exports` file

Remove all PowerHA SystemMirror managed exports from the AIX `/etc/exports` file for each PowerHA SystemMirror cluster node. Remove any file entries from the PowerHA SystemMirror `/usr/es/sbin/cluster/etc/exports` file that are not exported from PowerHA SystemMirror for each PowerHA SystemMirror cluster node.

To remove all the entries in the `/etc/exports` file on each PowerHA SystemMirror cluster node, run the following command:

```
cat /dev/null > /etc/exports
```

Upgrading a PowerHA SystemMirror cluster

These topics provide instructions for upgrading an existing PowerHA SystemMirror cluster configuration.

Migrating from PowerHA SystemMirror 6.1 to PowerHA SystemMirror 7.1.3 or PowerHA SystemMirror 7.2.0

To migrate from PowerHA SystemMirror Version 6.1 to PowerHA SystemMirror 7.1.3 or PowerHA SystemMirror 7.2.0, you must update the version of the AIX operating system for nodes in the cluster.

Note: You cannot upgrade from PowerHA SystemMirror Version 6.1 to PowerHA SystemMirror Version 7.2.1. If you are running PowerHA SystemMirror Version 6.1 and want to upgrade to PowerHA SystemMirror Version 7.2.1, you must first upgrade to PowerHA SystemMirror Version 7.1.3 or PowerHA SystemMirror Version 7.2.0.

PowerHA SystemMirror 7.1.0, or later, uses Cluster Aware AIX (CAA), which is a different clustering technology than what is used in PowerHA SystemMirror 6.1.

You can use the PowerHA SystemMirror 6.1 migration assistant tool called `clmigcheck` to start the migration evaluation process. This tool provides an interactive interface to collect information about your cluster and environment. The `clmigcheck` tool evaluates all the information that it collects and checks for possible incompatibility issues for the migration.

Planning for migrating from PowerHA SystemMirror 6.1 to PowerHA SystemMirror 7.1.3 or PowerHA SystemMirror 7.2.0

Before you start the migration process, you must plan and prepare your environment for the migration.

Software requirements

It is recommended that before you start the migration process, that you update your environment to PowerHA SystemMirror 6.1 with Service Pack 15.

Note: You cannot upgrade from PowerHA SystemMirror Version 6.1 to PowerHA SystemMirror Version 7.2.1. If you are running PowerHA SystemMirror Version 6.1 and want to upgrade to PowerHA SystemMirror Version 7.2.1, you must first upgrade to PowerHA SystemMirror Version 7.1.3 or PowerHA SystemMirror Version 7.2.0.

The following website lists the versions of PowerHA SystemMirror and the versions of the AIX operating system that are supported for the migration process: https://aix.software.ibm.com/aix/ifixes/PHA_Migration/ha_install_mig_fixes.htm

For the migration process to complete successfully, you must apply a few IFIX bundles to PowerHA SystemMirror and to the AIX operating system. These IFIX bundles include migration fixes and a set of fixes for the migration assistant tool called **clmigcheck**. The **clmigcheck** tool is updated with more verification checks on your environment. If the **clmigcheck** tool finds problems with your environment, the migration process fails and displays messages on why the failure occurred and how to fix the problem. The IFIX bundles also adds the **-V** option to the **clmigcheck** tool. You can use the **-V** option to walk through the migration process without doing the migration on your environment. In other words, you can use the **-V** option to complete a practice migration. The **-V** option also provides you with a list of issues that you might need to fix before you start the migration process.

You can download the IFIX bundles for the AIX operating system and PowerHA SystemMirror from the following website: https://aix.software.ibm.com/aix/ifixes/PHA_Migration/ha_install_mig_fixes.htm

Note: The IFIX bundles contain packages for Cluster Aware AIX (CAA), Reliable Scalable Cluster Technology (RSCT), and PowerHA SystemMirror. Before you install these packages, review the following information:

- The CAA and RSCT packages must be installed before you migrate to a newer version of PowerHA SystemMirror. You must reboot your system after you install the CAA and RSCT packages. Therefore, it is recommended that you install both the CAA and RSCT packages at the same time and then reboot the system by using the **shutdown -r** command.
- Review the readme file instructions that are included in the PowerHA SystemMirror package. If the package includes a fix to the cluster manager, you must stop cluster services before you install this package. To avoid any downtime, you can install the package immediately after you migrate to the latest version of PowerHA SystemMirror.

Migration and cluster environment checks

Review the following information before you start the migration process:

- Verify that the network, storage, and other components do not have any issues. You can plan for migration only when you have a stable PowerHA SystemMirror 6.1 cluster. To check the status of the cluster, run the **clstat** command.
- Create a backup of the cluster configuration by using the PowerHA SystemMirror snapshot function.
- Verify that each cluster node has a valid PowerHA SystemMirror license. If you have questions about your PowerHA SystemMirror license, contact your IBM support representative.
- Verify that you have the correct system privileges (root user) to install the software requirements as part of the migration process.

Invalid configurations for migration

The following configurations cannot be updated to from PowerHA SystemMirror 6.1:

- Configurations with FDDI, ATM, X.25, and token ring, cannot be migrated and must be removed from the configuration.
- Configurations with IPAT via replacement or hardware address takeover cannot be migrated and must be removed from the configuration.

- Configurations with heartbeat via aliasing cannot be migrated and must be removed from the configuration.
- Configurations with non-IP networks, such as RS 232, TMSCSI, TMSSA, and disk heartbeat cannot be configured in PowerHA SystemMirror Version 7.1, or later. A non-IP network configuration cannot be migrated. If you attempt to migrate a non-IP network configuration, it is removed during the migration processes.
- Logical Volume Manager (LVM) split-site configurations with disks that are assigned to each site in an PowerHA SystemMirror 6.1 cluster cannot be migrated. To achieve a similar configuration in PowerHA SystemMirror 7.1.3, or later, you can use SMIT and C-SPOC to configure mirror pools that belong to a volume group for LVM split-site mirroring.

Prerequisites for migrating from PowerHA SystemMirror 6.1 to PowerHA SystemMirror 7.1.3 or PowerHA SystemMirror 7.2.0

After you reviewed the planning information, you must complete prerequisite steps to prepare for migration to PowerHA SystemMirror 7.1.3 or PowerHA SystemMirror 7.2.0 that uses Cluster Aware AIX (CAA).

Note: You cannot upgrade from PowerHA SystemMirror Version 6.1 to PowerHA SystemMirror Version 7.2.1. If you are running PowerHA SystemMirror Version 6.1 and want to upgrade to PowerHA SystemMirror Version 7.2.1, you must first upgrade to PowerHA SystemMirror Version 7.1.3 or PowerHA SystemMirror Version 7.2.0.

To prepare your cluster for migration from PowerHA SystemMirror 6.1, complete the following steps:

Note: If you do not complete the following steps, you might create a partial cluster that will not function correctly after the migration process.

1. After you install an updated version of the AIX operating system and restart the system, you must verify that CAA related entries were successfully added in the `/etc/inetd.conf` file, the `/etc/services` file, and the `/etc/inittab.conf` file. To verify that CAA entries are added to all of files, run the following command:

```
# egrep "caa|clusterconf" /etc/services /etc/inetd.conf /etc/inittab
```

The command displays the output information similar to the following example:

```
/etc/services:clcomd_caa      16191/tcp
/etc/services:caa_cfg      6181/tcp
/etc/inetd.conf:caa_cfg stream tcp6 nowait root /usr/sbin/clusterconf clusterconf >>/var/adm/ras/clusterconf.log 2>&1
/etc/inittab:clusterconf:23456789:once:/usr/sbin/clusterconf
```

2. CAA uses host names that are defined by TCP/IP settings to identify whether a node is part of a cluster. Therefore, all host names must have unique names. Verify the following information about host names before you start the migration process:
 - a. If you use a service label as the nodes host name, you must set the host name to the actual persistent host name of the LPAR during the migration process. After the migration is complete, you can reset the host name to a service label by using the **hostname** command.
 - b. Verify that the `COMMUNICATION_PATH` variable for PowerHA SystemMirror is set to the same value as the persistent host name of the LPAR. The value for the persistent host name is defined as part of the TCP/IP inet0 definition. To check the value of the TCP/IP inet0 definition, run the **lsattr -El inet0** command.
3. Verify that each node in the cluster has an updated `/etc/hosts` file that contains all cluster member-related entries. Verify for each node in the cluster, that the host name in the `/etc/hosts` file follows the format such that, the FQDN resolvable long name is next to the IP address. In the following example for two nodes, the first node is named node111, IP address is 1.1.1.1, and the FQDN long name is node111.xxx.yyy.com:

```
1.1.1.1      node111.xxx.yyy.com      node111
1.1.1.2      node222.xxx.yyy.com      node222
```


4. Verify that the only line that is not commented out in the `/etc/netshvc.conf` file is `hosts = local4,bind4`.
5. Verify that the `/etc/cluster/rhosts` file on all nodes includes the FQDN long name.
6. Verify that the `inetd` daemon is running on all nodes.
7. Verify that the disk that is used as the repository disk is set to `reserve_policy = no_reserve` on all nodes. To verify that the disk has the correct setting, run the `lsattr -El <hdiskX> | grep reserve_policy` command, where `<hdiskX>` is the name of the disk. To set the value of the repository disk to `reserve_policy = no_reserve`, run the `chdev -l <hdiskX> -a reserve_policy=no_reserve` command, where `<hdiskX>` is the name of the disk.

Note: EMC PowerPath Version 5.5, or earlier, uses a different ODM attribute (`reservation_lock`) to manage SCSI reserves. You must reset the EMC reserve.

8. Verify that the disk that is used as the repository disk has the same PVID across all nodes in the cluster. To check that the PVID is the same across all nodes, run the `lspv` command.

Upgrading to PowerHA SystemMirror Version 7.2.1 prerequisites

You can upgrade to PowerHA SystemMirror Version 7.2.1 by using rolling migration, snapshot conversion, or offline migration.

Before you upgrade to PowerHA SystemMirror, you must have a basic knowledge of the following information:

- PowerHA SystemMirror (from high-level concepts to low-level tasks, such as planning, maintenance, and troubleshooting), because the upgrading process builds on that knowledge.
- You can only upgrade to PowerHA SystemMirror Version 7.2.1 from PowerHA SystemMirror 7.1.3 or PowerHA SystemMirror 7.2.0.

Note: You cannot upgrade from PowerHA SystemMirror Version 6.1 to PowerHA SystemMirror Version 7.2.1. If you are running PowerHA SystemMirror Version 6.1 and want to upgrade to PowerHA SystemMirror Version 7.2.1, you must first upgrade to PowerHA SystemMirror Version 7.1.3 or PowerHA SystemMirror Version 7.2.0.

- PowerHA SystemMirror Version 7.2.1 supports a non-disruptive upgrade. For more information, see the “Performing a non-disruptive upgrade from PowerHA SystemMirror 7.1.3 to PowerHA SystemMirror 7.2.0, or later” on page 15 topic.

The following table shows the version of PowerHA SystemMirror and the minimum required technology level of the AIX operating system.

AIX operating system requirements

To upgrade to PowerHA SystemMirror Version 7.2.1, your system must be running one of the following version of the AIX operating system:

- IBM® AIX 7 with Technology Level 3, or later
- AIX Version 7.2, or later

Note: You can install the latest available fixes for the AIX operating system from the Fix Central website.

Host name requirements

Review the following information about host name requirements when upgrading to PowerHA SystemMirror Version 7.2.1:

- The host name, the Cluster Aware AIX (CAA) node name, and the name in the `COMMUNICATION_PATH` field of the HACMPnode Object Data Manager (ODM), must be the same.
- The host name and the PowerHA SystemMirror node name can be different.

- The host name can be changed after the cluster is deployed in an environment.
- The following statements do not apply if the host name is configured by using the **hostname** command:
 - The host name cannot be a service address.
 - The host name cannot be an IP address that is on a network that is defined as private in PowerHA SystemMirror.

Host name requirements

Review the following information about host name requirements when upgrading to PowerHA SystemMirror Version 7.2.1:

- The host name, the Cluster Aware AIX (CAA) node name, and the name in the **COMMUNICATION_PATH** field of the HACMPnode Object Data Manager (ODM), must be the same.
- The host name and the PowerHA SystemMirror node name can be different.
- The host name can be changed after the cluster is deployed in an environment.
- The following statements do not apply if the host name is configured by using the **hostname** command:
 - The host name cannot be a service address.
 - The host name cannot be an IP address that is on a network that is defined as private in PowerHA SystemMirror.

Split and merge policy requirements

If the cluster already has split and merge policies configured, then these policies are converted during the upgrade process to the policies that are supported in PowerHA SystemMirror Version 7.2.1. The following table displays the split and merge policy combinations before you upgrade and after you upgrade to PowerHA SystemMirror Version 7.2.1:

Table 1. Split and merge policy conversion

Before migration: split and merge policies	After migration: split and merge policies
None - Majority	None - Majority
None - Priority	None - Majority
None - Manual	None - Majority
Tie breaker - Tie breaker	Tie breaker - Tie breaker
Tie breaker - Priority	Tie breaker - Tie breaker

Planning for an upgrade

Before you can start doing an update, you must plan correctly for the upgrade.

The following configurations cannot be updated to PowerHA SystemMirror Version 7.2, or later.

- Configurations with FDDI, ATM, X.25, and token ring, cannot be migrated and must be removed from the configuration.
- Configurations with IPAT via IP replacement or hardware address takeover cannot be migrated and must be removed from the configuration.
- Configurations with heartbeat via aliasing cannot be migrated and must be removed from the configuration.
- Configurations with non-IP networks, such as RS 232, TMSCSI, TMSSA, and disk heartbeat cannot be configured in PowerHA SystemMirror Version 7.2, or later. A non-IP network configuration cannot be migrated. If you attempt to migrate a non-IP network configuration, it is removed during the migration processes.

- Configurations with disks that are assigned to each site in an HACMP 6.1 cluster cannot be migrated. To achieve a similar configuration in PowerHA SystemMirror Version 7.2, or later, you can use SMT and C-SPOC to configure mirror pools that belong to a volume group for LVM split-site mirroring.

Planning is required to update to PowerHA SystemMirror Version 7.2, or later, because of the following technology:

- The repository is stored on a disk that must be SAN attached and zoned to be shared by every node in the cluster and only the nodes in the cluster.
- If you intend to use multicast communications for your cluster, you must use a multicast IP address to monitor the cluster.
- The product will assign multicast addresses, but you can explicitly specify multicast addresses.
- Ensure that multicast communication is functional in your network topology before migration.

Updating the current cluster configuration

You can upgrade your current cluster configuration.

To update the current cluster configuration, complete the following steps:

1. Check that all nodes in the cluster are up and running the same and most recent version of the PowerHA SystemMirror software. Check the IBM website for the latest PowerHA SystemMirror APARs and PTFs available for the current version.

To check cluster conditions and settings, complete the following steps:

- a. Use **clstat** to review the cluster state and to make certain that the cluster is in a stable state.
 - b. Review the **/etc/hosts** file on each node to make certain it is correct.
 - c. After AIX Version 6.1.6, or later is installed, enter the fully qualified host name of every node in the cluster in the **/etc/cluster/rhosts** file. You must do this on each node in the cluster.
 - d. Take a snapshot of a node in the cluster configuration.
 - e. Ensure that each cluster node has its own PowerHA SystemMirror license. Otherwise, contact an IBM representative about licensing PowerHA SystemMirror.
 - f. Ensure that you have privileges to perform the installation as the root user, or ask your system administrator to make the appropriate changes.
2. Review the installation prerequisites and ensure that the system being upgraded meets these requirements.
 3. If needed, upgrade the AIX operating system and RSCT before upgrading PowerHA SystemMirror. Doing a migration installation of AIX leaves the existing licensed programs (LP) installed and eliminate the need to reinstall a previous version of PowerHA SystemMirror.

Related information:

Monitoring a PowerHA SystemMirror cluster

Checking types of networks

Make sure that the version of PowerHA SystemMirror that you plan to upgrade to supports the types of networks that you plan to use. Remove or change unsupported types before you upgrade the PowerHA SystemMirror software.

If your previous configuration includes unsupported network types and you attempt to upgrade a node, the installation will fail and an error message will notify you to change the unsupported network type.

Supported networks include:

- Ethernet
- EtherChannel
- Virtual Ethernet
- Host Ethernet adapter, also known as Integrated Virtual Ethernet adapter.

Migration and security

PowerHA SystemMirror configuration database security changes might affect migration.

The PowerHA SystemMirror Configuration Database (ODM) has the following security enhancements:

- **Ownership.** All PowerHA SystemMirror ODM files are owned by the root user and the **hacmp** group. In addition, all PowerHA SystemMirror binary that are intended for use by non-root users are owned by root user and the **hacmp** group.
- **Permissions.** The **hacmpdisksubsystem** file is set with 600 permissions. Most of the other PowerHA SystemMirror ODM files are set with 640 permissions (the root user can read and write, while the **hacmp** group can only read). All PowerHA SystemMirror binary are intended for use by non-root users are installed with 2555 permissions (readable and executable by all users, with the **setgid** bit turned on so that the program runs as **hacmp** group).

During the installation, PowerHA SystemMirror creates the **hacmp** group on all nodes. By default, the **hacmp** group has permission to read the PowerHA SystemMirror ODMs, but does not have any other special authority. For security reasons, do not to expand the authority of the **hacmp** group.

If you use programs that access the PowerHA SystemMirror ODMs directly, you might need to rewrite them if they are intended to be run by non-root users. For example, all access to the ODM data by non-root users should be handled via the provided PowerHA SystemMirror utilities.

Before installing PowerHA SystemMirror, include the **hacmp** group in the master **/etc/group** file and propagate this change to all cluster nodes to prevent overwriting your **hacmp** group.

Priority override location and persistence

If you are upgrading from a previous release, PowerHA SystemMirror handles priority override location (POL) and persistence differently from previous releases.

These differences include:

- The priority override location (POL) setting is not used.
The POL setting is not used for resource groups that you move from one node to another. In general, if you move a resource group from one node to another, it remains on its new node until you need to move it again.
 - If you reboot the cluster (which you seldom need to do in PowerHA SystemMirror 5.4.1 and later), the group returns to the node that is originally configured as its highest priority node in the nodelist (if the group has a fallback policy that tells it to fall back).
 - If you do not reboot the cluster, the group remains on the node to which you moved it, and, if it has a fallback policy to fall back, then it falls back to its "acting" highest priority node.
- The persistence after a reboot is not retained.
If a resource group in your cluster has a fallback policy with the option **persist** across cluster reboot and resides on the node to which you moved it before an upgrade, when you upgrade to PowerHA SystemMirror 5.4.1, or later, then the resource group remains on its destination node after the upgrade. In this case, you did not reboot.
However, if you reboot the cluster, the group returns to the node that is its originally configured highest priority node.

Note: If you want the group to be permanently hosted on its originally configured highest priority node, change the highest priority node in the node list for the resource group.

Upgrade do's and don'ts

This section lists the major tasks to do during or before an upgrade, and it also lists what not to do.

Upgrade do's

Make sure you verify the following:

- Take a cluster snapshot and save it to the **/tmp** directory as well as to another machine and CD.
- Save a copy of any event script files to the **/tmp** directory as well as to another machine and CD.
- Ensure that the same level of cluster software (including PTFs) are on all nodes before beginning a migration.
- Ensure that the cluster software is committed (and not just applied).
- Run cluster verification and make sure no errors are reported. Then you will know your configuration is in a working state before you start the upgrade.
- During migration from one version of PowerHA SystemMirror to another version you can start and stop cluster services but you cannot use the manual option when starting cluster services. Manual startup is disabled when migration begins and cannot be used again until the migration is completed.

Note: In a parent and child resource group configuration with different home nodes, if you start cluster services on the node where the parent resource group is in an unmanaged state the corresponding child resource group is not released and reacquired.

Upgrade don'ts

During any type of upgrade, do not do the following tasks:

- Do not save your cluster configuration or customized event scripts under these directory paths: **/usr/es/sbin/cluster** or **/usr/lpp/cluster**. Data in these directories might be lost during an upgrade. Instead, copy files to a separate system or to a CD.
- Do not verify and synchronize the cluster configuration.
- Do not attempt a dynamic automatic reconfiguration (DARE). For example, do not change node priority in a resource group, add or remove a network, update an LVM component, or add users to a cluster.
- Do not use mixed versions of the cluster software for an extended period of time (hybrid state).

When migrating an active cluster, one node at a time (a rolling migration), the use of commands and functions are restricted as follows when the cluster has mixed versions (that is, it is in a hybrid state):

- Do not change the cluster topology or configuration.
- Do not run the **clrgmove** command from the command line or the SMIT interface to move resource groups during a migration.
- Do not use any **System Management (C-SPOC)** functions except to start or stop cluster services, or to move a resource group.
- Do not use the **Problem Determination Tools > View Current State** function.
- Do not use the **Cluster Nodes and Networks > Manage the Cluster > Snapshot Configuration** option or run the **clsnapshot** command.
- Do not use the **Problem Determination Tools > Recover From PowerHA SystemMirror Script Failure** option, or run the **clruncmd** command, except when running the command or SMIT option from the target node specified by the command.

Checking that the software is committed

Before upgrading the cluster, ensure that the current software installation is committed (not just applied).

To ensure that the software is already committed:

1. Run the **lslpp -h cluster.*** command.
2. If the word **APPLY** displays under the action header, enter **smit install_commit** before installing the PowerHA SystemMirror software.

SMIT displays the Commit Applied Software Updates (Remove Saved Files) panel.

3. Enter the following field values:

Table 2. Commit Applied Software Updates fields

Field	Value
SOFTWARE name	From the picklist, select all cluster file sets.
COMMIT old version if above version used it?	Set this field to Yes .
EXTEND filesystem if space needed?	Set this field to Yes .

Performing a rolling migration

You can upgrade a PowerHA SystemMirror cluster while keeping your applications running and available. During the upgrade process, a new version of the software is installed on each cluster node while the remaining nodes continue to run the earlier version. In this hybrid state, PowerHA SystemMirror continues to respond to cluster events.

Configuration changes are not allowed and new features cannot be activated until all nodes are migrated to the new release.

Performing a rolling migration from PowerHA SystemMirror 7.1.3, or later, to PowerHA SystemMirror 7.2.0, or later

To upgrade from PowerHA SystemMirror 7.1.3, or later, your system must be running IBM AIX 7 with Technology Level 3, or later, or AIX Version 7.2, or later.

You can migrate from PowerHA SystemMirror 7.1.3, or later, to PowerHA SystemMirror 7.2.0, or later, while keeping your applications up and running. During the migration, the new version of PowerHA SystemMirror is installed on each node in the cluster while the remaining nodes continue to run the earlier version of PowerHA SystemMirror. When your cluster is in this hybrid state, PowerHA SystemMirror still responds to cluster events. Until all nodes are migrated to the new version of PowerHA SystemMirror, you cannot make configuration changes and new functions are not active.

You can use the following steps to perform a rolling migration from your current version of PowerHA SystemMirror to the applicable upgraded version shown in the following table:

Table 3. PowerHA SystemMirror versions for performing a rolling migration

Current version of PowerHA SystemMirror	Upgraded version of PowerHA SystemMirror
PowerHA SystemMirror Version 7.1.3	PowerHA SystemMirror Version 7.2.0
PowerHA SystemMirror Version 7.1.3	PowerHA SystemMirror Version 7.2.1
PowerHA SystemMirror Version 7.2.0	PowerHA SystemMirror Version 7.2.1

To perform a rolling migration from, complete the following steps:

1. By using the **Move Resource Groups** option in System Management Interface Tool (SMIT), stop cluster services on a node that you want to migrate.
2. Install IBM AIX 7 with Technology Level 3, or later, or AIX Version 7.2, or later, on the node. When you install a newer version of the AIX operating system, a new version of Reliable Scalable Cluster Technology (RSCT) is also installed. Verify that the correct version of the AIX operating system and RSCT are working on the node.
3. Reboot the node by typing `shutdown -Fr`. When the node comes back online, it is running the new version of the AIX operating system and RSCT.
4. Install PowerHA SystemMirror 7.2.0, or later on the node. Verify that you are using a supported technology level of the AIX operating system for the version of PowerHA SystemMirror that you install.

- | **Note:** You must first install the base PowerHA SystemMirror filesets for the new release you are migrating to and then install any corresponding PowerHA SystemMirror service packs. Do not mix the base filesets for the new PowerHA SystemMirror release and the service packs in the same installation directory because it might affect the order of installation and cause errors.
- | 5. Using SMIT, start cluster services.
- | 6. Verify that the node is available in the cluster by typing `clmgr query cluster | grep STATE`.
- | 7. Repeat steps 1 - 7 on one node at a time for each node in the cluster.

| **Note:** You must bring cluster services online on all nodes in the cluster to complete the migration process.

| **Related reference:**

- | “Upgrading to PowerHA SystemMirror Version 7.2.1 prerequisites” on page 9
- | You can upgrade to PowerHA SystemMirror Version 7.2.1 by using rolling migration, snapshot conversion, or offline migration.
- | “Required versions of AIX and Reliable Scalable Cluster Technology (RSCT)” on page 23
- | PowerHA SystemMirror requires specific versions of AIX and RSCT and requirements for upgrading the AIX operating system.

| **Related information:**

- | Starting cluster services
- | Stopping cluster services

Performing a non-disruptive upgrade from PowerHA SystemMirror 7.1.3 to PowerHA SystemMirror 7.2.0, or later

You can use the Non-Disruptive Upgrade (NDU) function to update PowerHA SystemMirror software to a later version without any interruptions to resource groups and applications. The NDU function is available only if the PowerHA SystemMirror software upgrade does not require the AIX operating system to restart.

When you use the NDU function to upgrade PowerHA SystemMirror, the PowerHA SystemMirror software is updated on one node at a time in a cluster while the resource groups remain running. During the NDU process, the node is moved to an UNMANAGED state. This process allows the workloads that are managed by the resource groups remain functional while the upgrade occurs. After the NDU process completes, the node is moved back to a MANAGED state.

NDU migration is supported if you update from PowerHA SystemMirror Version 7.1.3 to PowerHA SystemMirror 7.2.0, or later, on a system that is running on one of the following versions of the AIX operating system:

- IBM AIX 6 with Technology Level 9
- IBM AIX 7 with Technology Level 3, or later
- IBM AIX Version 7.2, or later

The following are requirements and limitations for the NDU function:

- If your configuration is using the HyperSwap[®] function, you cannot use the NDU function to upgrade from PowerHA SystemMirror Version 7.1.3 to PowerHA SystemMirror 7.2.0, or later.
- If your configuration is using Geographic Logical Volume Manager (GLVM) with Remote Physical Volumes (RPV), you must unconfigure the RPV clients and servers before you use the NDU function.

To perform the NDU migration on a node, complete the following steps:

1. By using the **Unmanage Resource Groups** option in the System Management Interface Tool (SMIT), stop cluster services on the node that you want to upgrade the PowerHA SystemMirror software.

2. Install PowerHA SystemMirror 7.2.0, or later, on the node. Verify that you are using a supported technology level of the AIX operating system for the version of PowerHA SystemMirror that you want to install.

Note: You must first install the base PowerHA SystemMirror filesets for the new release you are migrating to and then install any corresponding PowerHA SystemMirror service packs. Do not mix the base filesets for the new PowerHA SystemMirror release and the service packs in the same installation directory because it might affect the order of installation and cause errors.

3. Using SMIT, start cluster services.
4. Repeat steps 1 - 3 on each node in the cluster, one node at a time.

Note: You must bring cluster services online on all nodes in the cluster to complete the migration process.

Upgrading PowerHA SystemMirror using a snapshot

You can upgrade a cluster by using a cluster snapshot that was created on the configuration with an earlier version of PowerHA SystemMirror. This upgrade path requires cluster services to be offline simultaneously on all nodes. As a result, your applications will not be highly available during the upgrade.

Upgrading from PowerHA SystemMirror 7.1.3, or later, by using a snapshot

To upgrade PowerHA SystemMirror by using a snapshot, you must already have a snapshot that is saved from PowerHA SystemMirror 7.1.3, or later.

You can use the following steps to upgrade PowerHA SystemMirror by using a snapshot from your current version of PowerHA SystemMirror to the applicable upgraded version shown in the following table:

Table 4. PowerHA SystemMirror versions for upgrading by using a snapshot

Current version of PowerHA SystemMirror	Upgraded version of PowerHA SystemMirror
PowerHA SystemMirror Version 7.1.3	PowerHA SystemMirror Version 7.2.0
PowerHA SystemMirror Version 7.1.3	PowerHA SystemMirror Version 7.2.1
PowerHA SystemMirror Version 7.2.0	PowerHA SystemMirror Version 7.2.1

To upgrade from PowerHA SystemMirror 7.1.3, or later, to PowerHA SystemMirror 7.2.0, or later, by using a snapshot, complete the following steps:

1. Using the System Management Interface Tool (SMIT), stop cluster services on all nodes in the cluster by using the **Bring resource groups offline** option.
2. Remove PowerHA SystemMirror from all nodes in the cluster.
3. Upgrade to a version of the AIX operating system that supports the version of PowerHA SystemMirror to which you plan to upgrade.
4. Install the version of PowerHA SystemMirror software to which you want to upgrade to on all the nodes in the cluster. Verify that you are using a supported technology level of the AIX operating system for the version of PowerHA SystemMirror that you want to install.
5. Restart all nodes in the cluster by running the **shutdown -Fr** command. When the nodes come back online, they will be running the new version of the AIX operating system and Reliable Scalable Cluster Technology (RSCT).
6. Convert the snapshot by running the following command, where *version* is the version of PowerHA SystemMirror (for example, 7.1.3 or 7.2.0) and *snapshotname* is the name of the snapshot that you are using for the upgrade.

```
/usr/es/sbin/cluster/conversion/clconvert_snapshot -v version -s snapshotname
```

7. Apply the converted snapshot by completing the following steps:

- a. From the command line, enter `smit sysmirror`.
- b. In SMIT, select **Cluster Nodes and Networks > Manage the cluster > Snapshot Configuration > Restore the Cluster Configuration from a Snapshot**, and press Enter.
- c. Select the snapshot that you want to restore, complete the required fields, and press Enter.

Note: The existing Cluster Aware AIX (CAA) cluster is removed when the snapshot is applied.

8. Run the `/usr/sbin/lsccluster -m` command to verify that the Cluster Aware AIX (CAA) services are active on each node in the cluster. If the CAA services are not available, contact IBM support to resolve the problem.
9. Using SMIT, start cluster services on the first node in the cluster.
10. Verify that each node is available in the cluster by typing `clmgr query cluster | grep STATE`.

Related tasks:

“Removing the PowerHA SystemMirror software” on page 30

Before you remove the PowerHA SystemMirror software from a system, stop cluster services. You cannot remove the software while the cluster is running.

Related reference:

“Upgrading to PowerHA SystemMirror Version 7.2.1 prerequisites” on page 9

You can upgrade to PowerHA SystemMirror Version 7.2.1 by using rolling migration, snapshot conversion, or offline migration.

Related information:

Saving and restoring cluster configurations

`clconvert_snapshot` command

Upgrading an offline cluster for PowerHA SystemMirror

You can take cluster services offline on all of the nodes and then update the cluster definitions individually on each node.

Related concepts:

“Upgrading PowerHA SystemMirror using a snapshot” on page 16

You can upgrade a cluster by using a cluster snapshot that was created on the configuration with an earlier version of PowerHA SystemMirror. This upgrade path requires cluster services to be offline simultaneously on all nodes. As a result, your applications will not be highly available during the upgrade.

Related reference:

“Performing a rolling migration” on page 14

You can upgrade a PowerHA SystemMirror cluster while keeping your applications running and available. During the upgrade process, a new version of the software is installed on each cluster node while the remaining nodes continue to run the earlier version. In this hybrid state, PowerHA SystemMirror continues to respond to cluster events.

Related information:

PowerHA SystemMirror use of Cluster Aware AIX

Upgrading an offline cluster from PowerHA SystemMirror 7.1.3, or later

The Cluster Aware AIX (CAA) function is able to upgrade clusters from PowerHA SystemMirror 7.1.3, or later, to PowerHA SystemMirror 7.2.0, or later. You do not have to remove any clusters when you are upgrading an offline cluster from PowerHA SystemMirror 7.1.3, or later.

You can use the following steps to upgrade an offline cluster from your current version of PowerHA SystemMirror to the applicable upgraded version shown in the following table:

Table 5. PowerHA SystemMirror versions for upgrading an offline cluster

Current version of PowerHA SystemMirror	Upgraded version of PowerHA SystemMirror
PowerHA SystemMirror Version 7.1.3	PowerHA SystemMirror Version 7.2.0
PowerHA SystemMirror Version 7.1.3	PowerHA SystemMirror Version 7.2.1
PowerHA SystemMirror Version 7.2.0	PowerHA SystemMirror Version 7.2.1

To upgrade an offline cluster from PowerHA SystemMirror 7.1.3, or later, to PowerHA SystemMirror 7.2.0, or later, complete the following steps:

1. Using the System Management Interface Tool (SMIT), stop cluster services on all nodes in the cluster that you want to upgrade.
2. Upgrade to a version of the AIX operating system that supports the version of PowerHA SystemMirror to which you plan to upgrade.
3. Restart all nodes in the cluster by running the **shutdown -Fr** command. When the nodes come back online, they will be running the new version of the AIX operating system and Reliable Scalable Cluster Technology (RSCT).
4. Upgrade to the version of PowerHA SystemMirror that you want to run on all the nodes in the cluster. Verify that you are using a supported technology level of the AIX operating system for the version of PowerHA SystemMirror that you are upgrading to.
5. Using SMIT, start cluster services on the first node in the cluster.
6. Verify that each node is available in the cluster by typing `clmgr query cluster | grep STATE`.

Related reference:

“Upgrading to PowerHA SystemMirror Version 7.2.1 prerequisites” on page 9

You can upgrade to PowerHA SystemMirror Version 7.2.1 by using rolling migration, snapshot conversion, or offline migration.

Related information:

Stopping cluster services

Starting cluster services

PowerHA SystemMirror use of Cluster Aware for AIX

Upgrade PowerHA SystemMirror using the `cl_езupdate` command

You can use the `cl_езupdate` command to update the software for the entire cluster or a subset of nodes in the cluster, often without disrupting workloads.

Some updates require a system reboot after the update is installed. If you install an update that requires a reboot, you must plan to move any active workloads to a different node that is not being updated.

You must run the `cl_езupdate` command from one of the nodes that is part of the cluster. When you run the `cl_езupdate` command, you must specify the location of the update images or an interim fix that you want to install. The location of the update images or interim fix can be either be a file system directory or a Network Installation Management (NIM) `lpp_source` resource that is located on a NIM server. The `cl_езupdate` command does not setup or configure the NIM environment. For more information about setting up NIM, see Network Installation Management.

You can complete the following tasks by using the `cl_езupdate` command:

- Query information about the cluster, nodes, NIM server, or service packs and interim fixes that are located in a specified installation location. The query can be run on the entire cluster or on a specific subset of nodes in the cluster.
- Apply and reject updates for AIX service packs or interim fixes. The `cl_езupdate` command cannot be used to update the cluster to newer AIX technology levels.

- Apply and reject updates for PowerHA SystemMirror service packs and technology levels, and interim fixes located in a specified installation location. This process is performed on the entire cluster or on a specific subset of nodes in the cluster. You can also apply updates in preview mode. When you use preview mode, all the prerequisites for installation process are checked, but the cluster updates are not installed on the system.
- Reject AIX service packs, PowerHA service packs, and interim fixes that were already installed on the system. This task is performed on the entire cluster or on a specific subset of cluster nodes.

Limitations

The following are limitations and restrictions for **cl_ezupdate** command:

- If you have previously installed any interim fixes, those fixes might be overwritten or removed when you apply a new service pack. If the previously installed interim fix has locked the fileset, you can override that lock and install the service pack by using the **-F** flag.
- You cannot install a new PowerHA SystemMirror technology level (TL) in the applied state. Filesets installed as part of new TL are automatically moved into committed state. This means that the installation image cannot be rejected. The **cl_ezupdate** tool cannot be used to uninstall technology levels.
- If you want to update the software by using a NIM resource, the NIM client must be configured first and must be available to all nodes where you want to use the **cl_ezupdate** tool.
- The **cl_ezupdate** tool requires an existing PowerHA SystemMirror and Cluster Aware AIX (CAA) cluster definition.
- The Cluster Communications daemon (**clcomd**) must be enabled to communicate with all nodes in the cluster. The **cl_ezupdate** tool attempts to verify **clcomd** communications before installing any updates.
- If a cluster node update operation fails, the **cl_ezupdate** script ends immediately and exits with an error. To troubleshoot the issue, an administrator must restart the update operation or undo the completed update operations.
- You must place any interim fixes in the **emgr/ppc** directory of the NIM **lpp_source** resource.
- The **cl_ezupdate** tool runs only on AIX version 7, or later.
- The **cl_ezupdate** tool can be used only on PowerHA SystemMirror 7.1.3, or later.
- The **cl_ezupdate** tool cannot be used with the AIX **multibos** utility.
- If you are running the **cl_ezupdate** tool on a cluster node that is not included as an option of the **-N** flag and if the **-S** flag specifies the file system path as an option, the cluster node on which you are running the command is the source node for install image propagation. This cluster node must have the file system path specified in the **-S** option.

Additional migration tasks

After you complete the upgrade, you might need to complete additional tasks.

Recompiling Clinfo clients after migrating

Recompiling existing Clinfo applications is not required.

However, in PowerHA SystemMirror, the **CL_MAXNAMELEN** value is 256 characters and there are changes related to resource group information.

The Clinfo interface works with Internet Protocol version 6 addresses.

If you want to incorporate these changes in your applications, make the desired modifications, recompile, and link the applications by using the Clinfo library.

Related information:

Programming client applications

Resetting PowerHA SystemMirror tunable values

In PowerHA SystemMirror, you can change the settings for a list of tunable values that were changed during the cluster maintenance and reset them to their default settings or to installation-time cluster settings.

The installation-time cluster settings are equivalent to the values that appear in the cluster after manually installing PowerHA SystemMirror.

Note: Resetting the tunable values does not change any other aspects of the configuration, while installing PowerHA SystemMirror removes all user-configured configuration information including nodes, networks, and resources.

List of tunable values

The following values can be reset:

- User-supplied information:
 - Cluster event customizations, such as all changes to cluster events. Resetting changes to cluster events does not remove any files or scripts that the customization used, only the PowerHA SystemMirror's knowledge of pre-event and post-event scripts.
 - Cluster event rules. Any changes made to the event rules database are reset to their installation-time default values.
 - PowerHA SystemMirror command customizations. Any changes to the default set of PowerHA SystemMirror commands are reset to their installation-time defaults.
- Automatically generated and discovered information:

Typically, you cannot see this information. PowerHA SystemMirror rediscovers or regenerates this information when the cluster services are restarted or during the next cluster synchronization.

PowerHA SystemMirror resets the following:

 - Local node names stored in the cluster definition database
 - Netmasks for all cluster networks
 - Netmasks, interface names, and aliases for heartbeat (if configured) for all cluster interfaces
 - Instance numbers and default log sizes for the RSCT subsystem

Related information:

Administering PowerHA SystemMirror

Troubleshooting PowerHA SystemMirror

Verifying the success of your migration

Now that all nodes have been migrated, you need to complete the necessary tasks to ensure that everything is working correctly.

Verifying software levels installed using AIX commands

Verify the software installation by using the `lppchk` AIX command, and check the installed directories to see that expected files are present.

The `lppchk` command verifies that files for an installable software product (file set) match the Software Vital Product Data (SWVPD) database information for file sizes, **checksum** values, or symbolic links.

Run the commands `lppchk -v` and `lppchk -c "cluster.*"`

If the installation is OK, both commands return nothing.

Automatically saved files

Several files in the `/usr/lpp/save.config` directory are automatically saved during the upgrade process.

These files include:

```
/usr/lpp/save.config/usr/es/sbin/cluster/events/node_up.rp
/usr/lpp/save.config/usr/es/sbin/cluster/events/node_down.rp
```

In addition, the following files are saved during the upgrade process and removed from the system at the end of migration:

```
/lpp/cluster/objrepos/PowerHA SystemMirror.npp
/lpp/cluster/objrepos/PowerHA SystemMirror.rude
```

CAUTION:

Until the cluster is migrated, do not delete any of the files listed above.

Verifying the upgraded cluster definition

After the upgraded PowerHA SystemMirror software is installed on all of the nodes in the cluster and cluster services are restored, verify and synchronize the cluster configuration. Verification ensures that the cluster definition is the same on all nodes. You can verify and synchronize a cluster only when all nodes in the cluster are running the same version of the software.

Before you verify the upgraded cluster definition, you must verify that the Cluster Aware AIX services are defined and active by running the following command:

```
/usr/sbin/lsccluster -m
```

To verify the cluster:

1. Enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Verify and Synchronize Cluster Configuration (Advanced) > Verify Changes only**, and press Enter.

Verifying all cluster file sets have migrated

Previous documentation APARs might not be successfully converted resulting in the inability to synchronize the cluster.

Execute the following to verify that all cluster file sets are at the expected level:

```
lspp -l | grep cluster
```

Running AIX commands on the migrated cluster

To determine which daemons are active on a cluster node, you can use the options from the following SMIT menu, **System Management (C-SPOC) > PowerHA SystemMirror Services > Show Cluster Services**

Execute the `/usr/es/sbin/cluster/utilities/clshowsrv -v` utility, which produces results similar to the following:

Status of the RSCT subsystems used by HACMP:

```
Subsystem Group  PID  Status
ctrmc  rsct    14458 active
```

Status of the HACMP subsystems:

```
SubsystemGroup  PID  Status
clcomd  clcomd  15778 active
clstrmgrEScluster  32792 active
```

Status of the optional HACMP subsystems:

```
SubsystemGroup  PID  Status
clcomd  cluster  31210 active
```

The `clcomd` daemon is now part of the AIX operating system. This requires that the fully qualified host names of all nodes in the cluster are listed in the `/etc/cluster/rhosts` file.

Troubleshooting your migration

Making configuration changes is not supported during the migration process. If you try to change the cluster topology or resources when migration is incomplete, the synchronization process fails and you get an error.

You receive a message similar to the following:

```
clclare: PowerHA SystemMirror migration is detected.  
DARE event cannot be run until the migration has completed.
```

When migration is complete, you can apply changes or remove them.

To remove changes, restore the active PowerHA SystemMirror configuration database:

1. Enter `smit sysmirror`.
2. In SMIT, select **Problem Determination Tools > Restore PowerHA SystemMirror Configuration Database from Active Configuration**.

Recovering your previous version of PowerHA SystemMirror

These topics discuss how to completely undo your migration and recover your previous version of PowerHA SystemMirror.

Recovering from a conversion failure

When you install PowerHA SystemMirror, the `cl_convert` command runs automatically to convert the PowerHA SystemMirror configuration database from a previous version of PowerHA SystemMirror to the current version. If the installation fails, run `cl_convert` to convert the database.

In a failed conversion, run `cl_convert` using the `-F` flag. For example, to convert from PowerHA SystemMirror 5.3, use the `-F` and `-v` (version) flags as follows:

```
cl_convert -F -v 5.3
```

To run a conversion utility, the following options are required:

- Root user privileges
- The PowerHA SystemMirror version from which you are converting

The `cl_convert` utility records the conversion progress to the `/tmp/clconvert.log` file so that you can gauge conversion success.

Recovering configuration information

The post-installation output informs you to merge site-specific configuration information into the newly installed files:

Some configuration files could not be automatically merged into the system during the installation. The previous versions of these files have been saved in a configuration directory as listed below. Compare the saved files and the newly installed files to determine if you need to recover configuration data. Consult your product documentation to determine how to merge the data.

Configuration files that were saved in `/usr/lpp/save.config`:

```
/usr/es/sbin/cluster/etc/rc.cluster  
/usr/es/sbin/cluster/samples/clinfo.rc  
/usr/es/sbin/cluster/samples/pager/sample.txt
```

```
/usr/es/sbin/cluster/etc/clinfo.rc<
/usr/es/sbin/cluster/utilities/clexit.rc
/usr/es/sbin/cluster/etc/clhosts
/usr/es/sbin/cluster/etc/rc.shutdown
/usr/es/sbin/cluster/diag/clconraid.dat
/usr/es/sbin/cluster/etc/hacmp.term
/etc/cluster/lunreset.lst
/etc/cluster/disktype.lst
```

Installing PowerHA SystemMirror on server nodes

These topics list the prerequisites for the PowerHA SystemMirror software and describe how to install it.

Related reference:

“Overview of the installation process” on page 1

After completing the planning steps, you are ready to install the PowerHA SystemMirror software.

Prerequisites for installing PowerHA SystemMirror on server nodes

Before you install the PowerHA SystemMirror software, ensure that your system meets the requirements listed in these topics.

Supported hardware

This topic discusses the hardware that is supported for PowerHA SystemMirror.

To ensure that your system meets the guidelines established for PowerHA SystemMirror, contact your sales representative, or see the IBM sales guide at the IBM Offering Information web site.

For a summary of hardware that is supported for PowerHA SystemMirror, see IBM Techdocs: PowerHA hardware support matrix (<http://www.ibm.com/support/techdocs/atmsastr.nsf/WebIndex/TD105638>).

Related concepts:

“OEM disk, volume group, and file systems in a cluster” on page 54

These topics describe how you can customize PowerHA SystemMirror software to integrate original equipment manufacturer (OEM) disks, volume groups, and file systems in a PowerHA SystemMirror cluster.

Related information:

 [IBM Offering Information](#)

Planning shared disk and tape devices

Disk space requirements

The PowerHA SystemMirror installation has certain disk space requirements.

These requirements include:

- The `/usr` directory must have 82 MB of free space for a full installation.
If you are not planning to install optional software, you can plan for less space. Also, install only the message catalogs for the language you will be using, rather than all message catalogs. (Japanese message catalogs use 1.6 MB.)
- The `/` (root) directory must have 710 KB of free space (beyond any need to extend the `/usr` directory).

Required versions of AIX and Reliable Scalable Cluster Technology (RSCT)

PowerHA SystemMirror requires specific versions of AIX and RSCT and requirements for upgrading the AIX operating system.

Requirements for upgrading AIX

Before you upgrade to a new version of the AIX operating system, review the following information:

1. Ensure that your adapter node numbers for all shared SCSI disk are not the same on each node. During an operating system upgrade, SCSI disk adapters are reset to the default value of 7. This setting can cause a conflict on the bus that prevents correct access to the shared disk.
2. Disable importation of volume groups. All volume groups are imported and varied on, on the node that is being upgraded. This causes loss of disk access by other nodes in the concurrent resource group in the cluster.
To prevent this disruption when migrating, use the option to disable importation of volume groups.
3. If you are planning to upgrade the AIX operating system at the same time as installing the PowerHA SystemMirror software, perform the AIX upgrade first.
If you are migrating from an environment using SNMP v.1, then before migrating PowerHA SystemMirror, execute the following series of commands:

```
stopsrc -s snmpd

/usr/sbin/snmpv3_ssw -l

startsrc -s snmpd
```
4. If you are migrating from an earlier version of AIX, the SNMP version might have been version 1. The default SNMP version is SNMP v3. You should migrate from and to the same version of SNMP to ensure that SNMP-based applications function correctly. After the migration has completed, you can switch to a different version of SNMP if you prefer. For example, if you are migrating from an environment that is using SNMP v.1, and you are upgrading to AIX 6.1, then before migrating PowerHA SystemMirror, execute the following series of commands: `stopsrc -s snmpd /usr/sbin/snmpv3_ssw -l startsrc -s snmpd`. SNMP Version 3 has some differences from Version 1.

Required AIX filesets

The following AIX filesets are required for PowerHA SystemMirror:

- bos.cluster.rte
- bos.ahafs
- bos.clvm.enh
- devices.common.IBM.storflow

Required RSCT filesets

The required versions of the RSCT filesets are installed automatically by default when you install AIX 6.1 with 6100-06, or later, or AIX Version 7.1, or later.

Install the RSCT filesets before installing PowerHA SystemMirror. Verify that each node has the same version of RSCT.

To determine if the appropriate filesets are installed and their levels, issue the following commands:

```
/usr/bin/lslpp -l rsct.core.rmc

/usr/bin/lslpp -l rsct.basic

/usr/bin/lslpp -l rsct.compat.basic.hacmp

/usr/bin/lslpp -l rsct.compat.clients.hacmp
```

If these filesets are not present, install the latest version of RSCT filesets.

Related tasks:

“Performing a rolling migration from PowerHA SystemMirror 7.1.3, or later, to PowerHA SystemMirror 7.2.0, or later” on page 14

To upgrade from PowerHA SystemMirror 7.1.3, or later, your system must be running IBM AIX 7 with Technology Level 3, or later, or AIX Version 7.2, or later.

Related information:

/etc/snmpd.conf file

Security file set requirements

If you plan to use message authentication or encryption for PowerHA SystemMirror communication between cluster nodes, the necessary file sets must be installed on each node.

These file sets include:

- **rsct.crypt.des** for data encryption with Data Encryption Standard (DES) message authentication
- **rsct.crypt.3des** for data encryption standard Triple DES message authentication
- **rsct.crypt.aes256** for data encryption with Advanced Encryption Standard (AES) message authentication

You can install these file sets from the AIX Expansion Pack CD-ROM.

English and Japanese message catalogs

US English and Japanese message catalogs are available. Set the LANG variable to an appropriate locale value so that SMIT help is available in the desired language.

Note that in SMIT, pressing F1 displays help information in the correct language only if the LANG variable is set correctly. PowerHA SystemMirror supports the following locales:

- en_US
- ja_JP

Also ensure that the correct base system locale is installed. To list the installed locales, type:

```
locale -a
```

To list the active locale, type:

```
locale
```

The active locale is determined by the LANG environment variable setting. If LANG=en_US, the locale will be en_US.

Ensure that the proper PowerHA SystemMirror message catalogs for the chosen language have been installed. To list the message catalogs, type:

```
lsipp -l "cluster.msg*"
```

Installation medium

The PowerHA SystemMirror software installation medium contains multiple file sets, some of which you install on all cluster nodes and clients. These file sets provide the services for cluster membership, system management, configuration integrity and control, failover, and recovery. It also includes cluster status and monitoring facilities for programmers and system administrators.

The PowerHA SystemMirror Application Plug-in software (on the PowerHA SystemMirror installation medium) contains the network service plug-in images. This plug-in file set provides example scripts to start and stop the network service scripts for Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), and printer services. Each network service has start and stop scripts bundled in a file set. These scripts are provided as examples that can be customized for your environment.

Several prerequisites must be completed before setup begins. A setup wizard is included in each file set to assist with the setup after installation.

The plug-in file sets are as follows:

- **cluster.es.plugins.dns.** This plug-in file set provides scripts to start and stop the DNS server process, provides a script to confirm that configuration files are present and stored in a shared file system, and scripts called by the monitoring functions of PowerHA SystemMirror that check on DNS server process life.
- **cluster.es.plugins.dhcp.** The DHCP Network Service Plug-in provides scripts to start and stop the DHCP server process, provides a script to ensure that configuration files are present and stored in a shared file system, and scripts called by the monitoring functions of PowerHA SystemMirror that check on DHCP server process life.
- **cluster.es.assist.printserver.** This plug-in file set provides scripts to start and stop the print server process, provides a script to ensure that configuration files are present and stored in a shared file system, and scripts called by the monitoring functions of PowerHA SystemMirror that check on print server process life.

PowerHA SystemMirror installable images

The organization of cluster images on the PowerHA SystemMirror media allows you to make individual or multiple image selections through SMIT when installing the PowerHA SystemMirror software.

The Reliable Scalable Cluster Technology (RSCT) images are prerequisites for PowerHA SystemMirror and are packaged with all versions of the AIX operating system.

The **cluster.es.client**, **cluster.es.server**, and **cluster.es.cspoc** images contain PowerHA SystemMirror runtime executable files and are required.

The following table displays all the PowerHA SystemMirror packages and a description of each package:

Table 6. PowerHA SystemMirror packages

Package	Description
cluster.adt.es	PowerHA SystemMirror Client Samples
cluster.doc.en_US.assist.db2.pdf	PowerHA SystemMirror Smart Assist documentation
cluster.doc.en_US.assist.oracle.pdf	PowerHA SystemMirror Smart Assist documentation
cluster.doc.en_US.assist.websphere.pdf	PowerHA SystemMirror Smart Assist documentation
cluster.doc.LOCALE.es (Locales include en_US, Ja_JP and ja_JP)	PowerHA SystemMirror general documentation
cluster.es.assist	PowerHA SystemMirror Smart Assist binaries
cluster.es.cfs	Cluster File System Support
cluster.es.client	PowerHA SystemMirror Client
cluster.es.cspoc	CSPOC
cluster.es.nfs	NFS Support
cluster.es.server	PowerHA SystemMirror Server Utilities
cluster.license	License
cluster.es.migcheck	Migration verification enablement
cluster.man.LOCALE.es (Locales include en_US, Ja_JP and ja_JP)	Man pages
cluster.man.LOCALE.assist (Locales include en_US, Ja_JP and ja_JP)	Man pages for Smart Assists
cluster.msg.LOCALE.es (Locales include en_US, Ja_JP and ja_JP)	General messages

PowerHA SystemMirror installation choices

Install the PowerHA SystemMirror software on each cluster node (server) and on any client machines that run the **clinfo** daemon.

You can install the software in the following ways:

- From a Network Installation Management (NIM) server

- From a hard disk to which the software has been copied
- Directly from the installation medium

Note that you must accept the license agreement as you install. Each node requires a PowerHA SystemMirror software license.

A user with root privileges must perform the installation.

Related reference:

“Installing PowerHA SystemMirror on client nodes” on page 31

These topics describe how to install and configure the PowerHA SystemMirror software on client systems. They also describe how to edit files and scripts related to the Cluster Information (Cinfo) program.

Installing from an installation server

To install the PowerHA SystemMirror software in a cluster environment, you can create a PowerHA SystemMirror installation server (containing all PowerHA SystemMirror software installable images) on one node and then load the images onto the remaining cluster nodes. By creating an installation server, you can load the PowerHA SystemMirror software onto other nodes faster from the server than from other media.

PowerHA SystemMirror supports the Network Installation Management (NIM) program including the Alternate Disk Migration option. For instructions on creating an installation server, see the *AIX Installation guide* or the *AIX Network Installation Management guide and reference*.

After installing the PowerHA SystemMirror software, verify the software installation. For information about verifying the software installation, see *Completing the installation*.

Related tasks:

“Completing the installation” on page 30

After you install the PowerHA SystemMirror software, verify the software installation.

Related information:

AIX installation and migration

Network installation management

Installing from a hard disk

To install the PowerHA SystemMirror software from your hard disk, first copy the software from the installation medium to the hard disk.

After installing the PowerHA SystemMirror software, verify the software installation. For information about verifying the software installation, see *Completing the installation*.

Related tasks:

“Completing the installation” on page 30

After you install the PowerHA SystemMirror software, verify the software installation.

Copying PowerHA SystemMirror software to hard disk

This topic describes copying PowerHA SystemMirror software to hard disk.

To copy the PowerHA SystemMirror software to your hard disk:

1. Place the PowerHA SystemMirror CD into the CD-ROM drive.
2. Enter `smit bffcreate`.
3. In the **INPUT device / directory for software** field, enter the name of the CD-ROM drive and press Enter.

Select the correct drive from the list and press Enter. That value is entered into the **INPUT device/directory** field as the valid input device.

4. Press Enter to display the Copy Software to Hard Disk for Future Installation panel.
5. Enter the following values for each field:

Table 7. Copy Software to Hard Disk for Future Installation fields

Field	Value
SOFTWARE name	Make a selection from the picklist, or select all to copy all server and client images.
DIRECTORY for storing software	Change the value to the storage directory accessed by all nodes using PowerHA SystemMirror.

6. Enter values for the other fields as appropriate for your site.
7. When you are satisfied with the entries, press Enter.
SMIT prompts you to confirm your selections.
8. Press Enter again to copy the software.

Installing PowerHA SystemMirror from the hard disk

After the PowerHA SystemMirror software has been copied to your system, you need to install the software.

Install the software by following the instructions in Installing from the installation medium. In the **INPUT device / directory for software field**, specify the location where you have copied the software.

Related tasks:

“Installing from the installation medium”

If you install the PowerHA SystemMirror software from the CD-ROM, you install the software directly onto each cluster node.

Installing from the installation medium

If you install the PowerHA SystemMirror software from the CD-ROM, you install the software directly onto each cluster node.

To install the PowerHA SystemMirror software:

1. Insert the CD into the CD-ROM drive and enter:
`smit install_all`
2. Enter the device name of the installation medium or install directory in the **INPUT device / directory for software** field and press Enter.
If you are unsure about the input device name or about the installation directory, select one from the picklist.
3. Enter the following values for each field.

Note: Press F1 (help) for any field. Use F4 to list the software before proceeding with the installation. This way you can install either the English or the Japanese message catalogs, and you can omit optional software if you want.

Table 8. INPUT device / directory for software fields

Field	Description
INPUT device / directory for software	This field shows the device or directory you specified earlier.
SOFTWARE to install	Select an option from the picklist, or enter all to install all server and client images. For a list of file sets, also see the section PowerHA SystemMirror installable images. Install the cluster.es required images (which contain the PowerHA SystemMirror runtime executables) on all servers: Make sure to install the level of RSCT required for your release of AIX. For more information, see "Required versions of AIX and Reliable Scalable Cluster Technology (RSCT)" on page 23.
PREVIEW only	If set to Yes , the preview option will check and ensure that installation prerequisites are met (for example, that required software is installed and sufficient disk space is available). When you are ready to perform the actual installation, set this field to No .
COMMIT software updates	This field applies only when installing software updates (PTFs).
SAVE replaced files	Use the default. This field applies only when installing software updates (PTFs). If you select No to Commit software updates you must select yes for this field.
AUTOMATICALLY install requisite software	Use the default. Set this field to No if the prerequisite software for PowerHA SystemMirror is already installed or if the OVERWRITE same or newer versions? field is set to yes ; otherwise, set this field to Yes to install required software.
EXTEND file systems if space needed	Select Yes if you have adequate hard disk space, or No if you have limited space.
OVERWRITE same or newer versions	Use the default. For normal new installations, leave this field set to No . Set it to yes if you are reinstalling. If you set this field to Yes , you must set the Automatically install requisite software field to No .
VERIFY install and check file sizes	Use the default. Select Yes if you want the system to perform some checks on the software you installed.
DETAILED output	Select Yes if you want a detailed log of all installation messages.
Process multiple volumes	Select this option if you want to enable the processing of multiple-volume CDs.
ACCEPT new license agreements	Select Yes for this item to proceed with the installation. If you select No , installation might stop with a warning that one or more file sets require software license agreements. You accept the license agreement only once for each node.
Preview new license agreements	Select Yes to view the text of the license agreements. The text appears in the current window in the language defined on your system.

4. When you are satisfied with the entries, press Enter.
SMIT prompts you to confirm your selections.
5. Press Enter again to install.
6. After the installation is completed, verify the installation as described in Completing the installation.

Be sure to read the PowerHA SystemMirror release_notes file in the `/usr/es/sbin/cluster/` directory for further instructions and the latest information on requirements or known issues.

Related tasks:

"Completing the installation" on page 30

After you install the PowerHA SystemMirror software, verify the software installation.

Related reference:

“PowerHA SystemMirror installable images” on page 26

The organization of cluster images on the PowerHA SystemMirror media allows you to make individual or multiple image selections through SMIT when installing the PowerHA SystemMirror software.

“Required versions of AIX and Reliable Scalable Cluster Technology (RSCT)” on page 23

PowerHA SystemMirror requires specific versions of AIX and RSCT and requirements for upgrading the AIX operating system.

Completing the installation

After you install the PowerHA SystemMirror software, verify the software installation.

1. Verify the software installation by using the AIX command `lppchk` and check the installed directories for the expected files.

The `lppchk` command verifies that files for an installable software product (file set) match the Software Vital Product Data (SWVPD) database information for file sizes, checksum values, or symbolic links.

2. Run the commands `lppchk -v` and `lppchk -c "cluster.*"`

If the installation is OK, both commands should return nothing.

3. Reboot each PowerHA SystemMirror cluster node and client.

Obtaining PowerHA SystemMirror APAR

An authorized program analysis report (APAR) contains an account of a problem caused by a suspected defect in a current, unaltered release of a program.

You can obtain a list of PowerHA SystemMirror APARs and updates for hardware as follows:

1. Go to IBM support website at <http://www.ibm.com/support/us/en/>.
2. Search for "PowerHA SystemMirror +APAR".
3. Sort the results by date, most recent first.

Troubleshooting during the installation

If you experience problems during the installation, the installation program automatically performs a cleanup process. If the cleanup does not perform automatically, you can perform a cleanup manually.

To perform a manual cleanup after an unsuccessful installation:

1. To display the Installation and Maintenance menu, enter `smit install`.
2. Select **Install and Update Software**.
3. Select **Clean Up After a Interrupted Installation**.
4. Review the SMIT output (or examine the `/smit.log` file or `/smit.log.< locale >` file) for the interruption's cause.
5. Fix any problems and repeat the installation process.

Removing the PowerHA SystemMirror software

Before you remove the PowerHA SystemMirror software from a system, stop cluster services. You cannot remove the software while the cluster is running.

To stop cluster services, enter the fast path `smitty clstop` and press Enter.

To remove the PowerHA SystemMirror software and your cluster configuration on cluster nodes and clients:

1. Enter `smit install_remove`.
2. Enter the following values for the fields.

Table 9. Removing PowerHA SystemMirror software

Field	Value
SOFTWARE name	Enter cluster* to remove all server and client software, or you can select options from the picklist. Your selections appear in this field.
REMOVE dependent software	Select No.
EXTEND filesystems if space needed	Select Yes.
DETAILED output	Select No.

3. When you are satisfied with the entries, press Enter.

SMIT prompts you to confirm your selections.

4. Press Enter again to remove the software.

Related information:

Starting and stopping cluster services

Installing PowerHA SystemMirror on client nodes

These topics describe how to install and configure the PowerHA SystemMirror software on client systems. They also describe how to edit files and scripts related to the Cluster Information (Clnfo) program.

Installing the PowerHA SystemMirror software on each client that runs the **clnfo** daemon enables the clients to receive messages about events and actions taken by the high availability software running in the cluster. The client can take predefined, automatic steps in response to some situations handled by the high availability software. The client can also print messages to inform users logged in to a client of the cluster state and thus make them aware of actions required to maintain connectivity.

Prerequisites

Before you install the PowerHA SystemMirror software on client systems:

- Make sure that the system has at least 2.6 MB available for the installation.
- Read the Release Notes for PowerHA SystemMirror in `/usr/es/sbin/cluster/release_notes` for additional information about installing the PowerHA SystemMirror software.
- Install PowerHA SystemMirror on the server nodes.

Related concepts:

“Installing PowerHA SystemMirror on server nodes” on page 23

These topics list the prerequisites for the PowerHA SystemMirror software and describe how to install it.

Related information:

Planning PowerHA SystemMirror

Step 1: Install the base system client images

This topic describes how to install the base high availability software on a client.

For a new installation, the `/usr` directory requires a minimum of 2.6 MB of available space.

To install the base high availability software on a client:

1. Place the PowerHA SystemMirror CD into the CD-ROM drive and enter:


```
smit install_selectable_all
```

 - If you are not sure of the name of the input device, select one from the picklist.
 - Select the correct drive and press Enter. That value is entered into the **INPUT device/directory** field as the valid input device.
2. Press Enter. SMIT refreshes the panel.

3. In the **SOFTWARE to install** field, use the picklist to select the client software modules associated with the following cluster images: **cluster.es.client**, **cluster.msg**, and **cluster.adt.es**

Note: If you select at least one client module associated with an installable image, all other required client modules are installed automatically.

4. Enter values for other fields as appropriate for your site.
5. Press Enter when you are satisfied with the entries.
SMIT prompts you to confirm your selections.
6. Press Enter again.
7. Read the PowerHA SystemMirror **release_notes** file in the **/usr/es/sbin/cluster/** directory for further instructions.
8. Copy the **/usr/es/sbin/cluster/etc/clhosts.client** file from the server to each client node, renaming it as **/usr/es/sbin/cluster/etc/clhosts**.
9. Optional: Edit the **clinfo.rc** script as described in Step 3: Editing the clinfo.rc script.
10. Reboot the client.
11. Repeat this procedure for each client system.

Related reference:

“Step 3: Edit the clinfo.rc script”

After installing the base client images and copying the **clhosts.client** file onto client nodes, you need to edit the **clinfo.rc** script.

Step 2: Copy the clhosts.client file onto client nodes

During the verification phase during initial cluster synchronization, PowerHA SystemMirror creates and populates a **clhosts.client** file on all PowerHA SystemMirror servers.

Copy this **/usr/es/sbin/cluster/etc/clhosts.client** file from a PowerHA SystemMirror server node to all client nodes, renaming it **clhosts**, removing the **.client** extension.

Step 3: Edit the clinfo.rc script

After installing the base client images and copying the **clhosts.client** file onto client nodes, you need to edit the **clinfo.rc** script.

The **/usr/es/sbin/cluster/etc/clinfo.rc** script updates the address resolution protocol (ARP) cache on the local node whenever an event occurs. A copy of the **/usr/es/sbin/cluster/etc/clinfo.rc** script must exist on each server node and client in the cluster in order for all ARP caches to be updated.

The PowerHA SystemMirror software is distributed with a template version of the **/usr/es/sbin/cluster/etc/clinfo.rc** script. You can use the script as distributed, modify it, or replace it with a custom script.

Note: If you are not using IPAT via IP replacement with alternate hardware addresses, the ARP function must remain in the **clinfo.rc** script.

The format of the **clinfo** call to **clinfo.rc** follows:

```
clinfo.rc {join,fail,swap} interface_name
```

The **clinfo** utility obtains information about the interfaces and their current state, and checks for changed states of interfaces:

- If a new state is UP, **clinfo** calls **clinfo.rc join interface_name**.
- If a new state is DOWN, **clinfo** calls **clinfo.rc fail interface_name**.
- If **clinfo** receives a **node_down_complete** event, it calls **clinfo.rc** with the fail parameter for each interface currently UP.

- If **clinfo** receives a **fail_network_complete** event, it calls **clinfo.rc** with the fail parameter for all associated interfaces.
- If **clinfo** receives a **swap_complete** event, it calls **clinfo.rc swap interface_name**.

If you have custom applications that use the Clinfo API and plan to use symmetric multiprocessors, you might need to modify your application. Afterwards recompile and link your application.

Related information:

Planning PowerHA SystemMirror

Programming client applications

Step 4: Update the ARP cache for clients not using clinfo

You need to update the ARP cache for clients that are not using the **Clinfo** API.

To update the ARP cache for clients not using the **Clinfo** API:

1. On clients that are not running **Clinfo**, you will need PowerHA SystemMirror to update the client's ARP cache by pinging the client from a cluster node when an IP address has moved.
2. `ping -c1 $host`
3. Reboot the clients.

Configuring installed hardware

These topics describe how to ensure that network interface cards (NICs), shared external disk devices, and shared tape drives are ready to support a PowerHA SystemMirror cluster.

Before you read this section, you should have already installed the devices following the instructions in the relevant AIX documentation. For the latest information about the hardware supported for use with the version of AIX that you are using, see the IBM website.

Note: For information about installing and configuring OEM disks, see OEM disk, volume group, and file systems accommodation.

Related concepts:

“OEM disk, volume group, and file systems in a cluster” on page 54

These topics describe how you can customize PowerHA SystemMirror software to integrate original equipment manufacturer (OEM) disks, volume groups, and file systems in a PowerHA SystemMirror cluster.

Configuring network Ethernet interface cards

When using the **smit mktcpip** fast path to define a NIC, the **ifconfig** command is used to configure the interface, and the host name field changes the default host name.

After the primary NIC is configured on a system, use **smit chinnet** (or **chdev**) to configure additional interfaces to avoid changing the system host name.

Related information:

Planning cluster network connectivity

Installing and configuring shared Fibre Channel tape drives

Complete the procedures in this section to install a shared IBM Fibre Channel tape drive.

Prerequisites

Be sure to install the appropriate Fibre Channel adapters. The installation procedures outlined in this section assume you have already installed these adapters. To install an adapter, if you have not done so,

follow the procedure outlined in the documentation you received with the unit.

The IBM Fibre Channel tape drive also requires the installation of IBM AIX Enhanced Tape and Medium Changer Device Drives (Atape). Follow the installation directions that are included with this device driver.

Logical device names

All nodes that are connected to the Fibre Channel tape drive must have the same logical device name (`/dev/rmt0`).

If they differ (for example, `/dev/rmt0` and `/dev/rmt1`) perform the steps in the following procedure:

To configure logical device names:

1. On the node with the lower numbers, put the tape device in the **Defined** state with the `rmdev -l` command. For example, run `rmdev -l rmt0`.
2. Enter `smit` tape.
3. From the Tape Drive panel, select **Change/Show Characteristics of a Tape Drive**.
4. Select the drive you put in the **Defined** state and note its characteristics as **Defined** in the various fields.
5. From the Tape Drive panel, select **Add a Tape Drive**.
6. Use the information gathered to create a new logical device name.
7. Review both nodes to assure that the logical device name is in the **Available** state and that the external tape drive has this same logical name.

Configuring the installation of a shared tape drive

During the boot process, the AIX operating system configures all the devices that are connected to the I/O bus, including the SCSI adapters. The AIX operating system assigns each adapter a logical name of the form `scsi` or `fscsi`, where *x* is an integer.

For example, an adapter could be named `scsi0` or `fscsi0` (Fast/Wide Adapters are named `scsix`). After the AIX operating system configures the SCSI adapter, it probes the SCSI bus and configures all target devices connected to the bus.

To confirm the installation of a tape drive:

1. Ensure that the AIX operating system created the device definitions that you expected.

To determine the logical name of the tape drive, use the `lsdev` command as follows:

```
lsdev -Cc tape
```

For example, the `lsdev` command might return the following output:

```
Name  Status Location Description
rmt0  Available00-02-00-000ther SCSI Tape Device
```

2. Ensure that the logical device name for your tape device has a status of **Available**.

If the tape drive has a status of **Defined** (instead of **Available**) ensure that the cable connections are secure, and then use the `mkdev` command to make the tape available. Enter:

```
mkdev -l rmtx
```

where *rmtx* is the logical name of the defined tape drive.

At this point, your tape drive installation is complete.

Related information:

Administering PowerHA SystemMirror

Defining shared LVM components

These topics describe how to define the LVM components shared by the nodes in a PowerHA SystemMirror cluster.

Creating the volume groups, logical volumes, and file systems shared by multiple nodes in a PowerHA SystemMirror cluster requires that you perform steps on all those nodes. In general, you define the components on one node (referred to in the text as the source node) and then import the volume group onto each of the other nodes that will be part of the resource group (referred to as destination nodes). This ensures that the definitions in the PowerHA SystemMirror configuration database for the shared components are the same on all the nodes that may access them.

Using SMIT, you can collect information about all volume groups available either on a local node, or on all nodes in the cluster defined to the resource group. Later, you can automatically import these volume groups onto all the destination nodes, if needed. As you configure shared resources in SMIT, when you add a volume group to a resource group, you can select the option to automatically import the volume group onto all the destination nodes in the participating resource group.

PowerHA SystemMirror nonconcurrent access environments typically use journaled file systems to manage data, while concurrent access environments use raw logical volumes. This topic provides instructions for both types of environments.

While configuring file systems for a resource group, you can select one or more individual file systems to be mounted in the particular resource group. In addition, you can specify that all file systems in the specified shared volume groups are mounted at run time. Using this option allows you to add or delete file systems in the particular shared volume group without resynchronizing the cluster resources after each update.

Defining shared LVM components for nonconcurrent access

PowerHA SystemMirror nonconcurrent access environments typically use journaled file systems to manage data, although some database applications may bypass the journaled file system and access the logical volume directly.

The key consideration is whether the environment uses mirrors. Shared logical volumes that reside on non-RAID disks must be mirrored in AIX to eliminate the disk as a single point of failure. Shared volume groups residing on IBM 2105 E10 and E20 Enterprise Storage Server® RAID devices and IBM TotalStorage DS6000™ and DS8000® storage systems in the AIX operating system. These devices provide their own data redundancy.

Defining shared LVM components

You can define shared LVM components with or without mirroring.

To define shared LVM components on the source node:

1. Create the volume group.
2. Create the journaled file system.
3. Rename the **jfslog** and the logical volume so that they are unique across the cluster.
4. Mirror the **jfslog** and logical volume (only with AIX mirroring on one node in the cluster).
5. Vary off the volume group.

To define shared LVM components (with or without AIX mirroring) on the other nodes that will access this volume group:

1. Import the volume group.
2. Change the volume group options so that it does not vary on at system startup.

- Vary off the volume group.

Creating a shared volume group on source node

You can create a shared volume group on a source node.

To create a shared volume group on a source node:

- Enter `smit mkvg`.
- Use the default field values unless your configuration has other requirements, or unless you are specifically instructed otherwise here. (Only fields with specific PowerHA SystemMirror instructions are shown.)

Note: If you are creating a mirrored volume group, select two or three times as many disks for the volume group as you will need for your data, depending on the number of mirrors you select.

Table 10. PowerHA SystemMirror volume group fields

Field	Value
VOLUME GROUP name	The name of the shared volume group must be unique within the cluster and distinct from the service IP label or IP address and resource group names. It should relate to the application it serves, as well as to any corresponding device, such as <code>websphere_service_address</code> .
Activate volume group AUTOMATICALLY at system restart	Set to No so that the volume group can be activated as appropriate by the cluster event scripts.
Volume Group MAJOR NUMBER	If you will be using NFS, make sure to use the same major number on all nodes. Use the <code>lvlstmajor</code> command on each node to determine a free major number common to all nodes.
Create VG Concurrent Capable?	This field must be set to enhanced concurrent for any shared volume group used in the cluster.

Creating a shared file system on source node

You can create a shared file system on a source node.

To create a shared file system on a source node:

- Enter `smit crjfs`.
When you create a journaled file system, AIX creates the corresponding logical volume. Therefore, you do not need to define a logical volume.
- Rename both the logical volume and the log logical volume for the file system and volume group.
- Review the settings for the following fields:

Table 11. Journaled file system fields

Field	Value
Mount AUTOMATICALLY at system restart	Make sure this field is set to No .
Start Disk Accounting	Set this field to No unless you want disk accounting.

Renaming jfslog logs and logical volumes on source node

AIX assigns a logical volume name to each logical volume it creates. Examples of logical volume names are `/dev/lv00` and `/dev/lv01`. Within a PowerHA SystemMirror cluster, the name of any shared logical volume must be unique. Also, the journaled file system log (**jfslog**) is a logical volume that requires a unique name in the cluster.

To make sure that logical volumes have unique names, rename the logical volume associated with the file system and the corresponding **jfslog** logical volume. Use a naming scheme that indicates that the logical volume is associated with a certain file system. For example, `lvsharefs` could name a logical volume for the `/sharefs` file system.

To rename **jfslogs** and logical volumes:

1. Use the **lsvg -l volume_group_name** command to determine the name of the logical volume and the log logical volume (**jfslog**) associated with the shared volume groups. In the command output:
 - Look for the logical volume name that has file type **jfs**. This is the logical volume.
 - Then look for the logical volume name that has file type **jfslog**. This is the log logical volume.
2. Use the **smit chl v** fast path to rename the logical volume and the log logical volume.
3. After renaming the **jfslog** or a logical volume:
 - Check the **/etc/filesystems** file to make sure that the **dev** and **log** attributes reflect the change.
 - Check the **log** attribute for each file system in the volume group and make sure that it has the new **jfslog**.
 - Check the **dev** attribute for the logical volume that you renamed and make sure that it has the new logical volume name.

Adding copies to logical volume on source node

You can add logical volume copies on a source node.

Note: This step is not needed for disk subsystems that provide their own mirroring, such as the IBM 2105 Enterprise Storage Server and IBM TotalStorage DS6000 and DS8000 storage systems.

To add logical volume copies on a source node:

1. Use the **smit mklv copy** fast path to add copies to a logical volume. Add copies to both the **jfslog** log logical volume and the logical volumes in the shared file systems. To avoid space problems, first mirror the **jfslog** log logical volume and then the shared logical volumes.

The copies should reside on separate disks that are controlled by different disk adapters and are located in separate drawers or units, if possible.

2. Review the number of logical volume copies. Enter:

```
lsvg -l volume_group_name
```

In the command output, locate the line for the logical volume for which you just added copies. Notice that the number in the physical partitions column is x times the number in the logical partitions column, where x is the number of copies.

3. To review the placement of logical volume copies, enter:

```
lspv -l hdiskx
```

where *hdiskx* is the name of each disk to which you assigned copies. That is, you enter this command for each disk. In the command output, locate the line for the logical volume for which you just added copies. For copies placed on separate disks, the numbers in the logical partitions column and the physical partitions column should be equal. Otherwise, the copies were placed on the same disk and the mirrored copies will not protect against disk failure.

Testing a file system

Run a consistency check on each file system's information.

To run a consistency check:

1. Enter:

```
fsck /filesystem_name
```

2. Ensure that you can mount the file system by entering:

```
mount /filesystem_name
```

3. Ensure that you can unmount the file system by entering:

```
umount /filesystem_name
```

Varying off a volume group on the source node

After testing the file system, use the **varyoffvg** command to deactivate the shared volume group.

You vary off the volume group so that it can be properly imported onto the destination node and activated as appropriate by the PowerHA SystemMirror event scripts. Enter the following command:
`varyoffvg volume_group_name`

Collecting information on current volume group configuration

PowerHA SystemMirror can collect information about all volume groups available across all nodes in the cluster as part of the discovery process.

PowerHA SystemMirror collects information about all shared volume groups that are currently available on the nodes in the cluster. It then determines which volume groups can be imported to the other nodes in the resource group. PowerHA SystemMirror filters out those volume groups that are already included in any other resource groups. You can use this information later to import discovered volume groups onto other nodes in the resource group that do not have these volume groups.

To collect the information about volume group configuration:

1. Enter `smit sysmirror`.
2. In SMIT, select **Cluster Nodes and Networks > Discover Network Interfaces and Disks** and press Enter. The information on current volume group configuration is collected as part of the general configuration information.

Importing a volume group onto destination nodes

Importing the volume group onto the destination nodes synchronizes the definition in the PowerHA SystemMirror configuration database for the volume group on each node on which it is imported.

When adding a volume group to the resource group, you can manually import a volume group onto the destination nodes or automatically import it onto all the destination nodes in the resource group.

Importing a volume group automatically:

You can set up the automatic import functions for volume groups.

For PowerHA SystemMirror to import available volume groups, ensure that the following conditions are met:

- Volume group names must be the same across cluster nodes, and unique to the cluster.
- Logical volumes and file systems must have unique names.
- All physical disks must be known to AIX and have PVIDs assigned.
- A volume group can be automatically imported to other nodes only if the physical disks on which it resides are available to all of the nodes in the resource group.

To automatically import a volume group:

1. Enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resource Groups > Change/Show Resources and Attributes for a Resource Group** and press Enter.

This function enables PowerHA SystemMirror to automatically import shareable volume groups to all the destination nodes in the resource group. With the automatic import function you can create a volume group and then add it to the resource group immediately, without manually importing it to each of the destination nodes in the resource group.

Before automatically importing volume groups, make sure that you have collected the information on appropriate volume groups by using the **Discover Network Interfaces and Disks** option.

Each volume group is assigned a major number when it is created. When PowerHA SystemMirror automatically imports a volume group, the major number already assigned to the volume group will be

used if it is available on all the destination nodes. Otherwise, any free major number will be used. NFS failover requires that the major number be the same on all nodes.

Related information:

Using NFS with PowerHA SystemMirror

Adding volume groups to resource groups automatically:

You can add a volume group to a resource group using automatic import.

To add volume groups to a resource group:

1. Enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resource Groups > Change/Show Resources and Attributes for a Resource Group** and press Enter.
3. Select the resource group for which you want to define a volume group and press Enter. A panel is shown with fields for all types of resources.
4. The **Automatically Import Volume Groups** flag is set to **False** by default. Set this flag to **True**.
5. In the **Volume Groups** field, select the volume groups from the picklist or enter volume groups names.

If you ran the discovery process before starting this procedure, review the list of all existing volume groups clusterwide (excluding **rootvg**). The list includes all shared volume groups in the resource group and the volume groups that are currently available for importation onto the resource group nodes. PowerHA SystemMirror filters out from the clusterwide list those volume groups that are already included in any of the resource groups.

6. Press Enter. PowerHA SystemMirror determines whether the volume groups that you entered or selected in the **Volume Groups** field need to be imported to any of the nodes that you defined for the resource group, and proceeds to import them, as needed.

Final state of automatically imported volume groups:

When PowerHA SystemMirror automatically imports volume groups, the final state (varied on or varied off) depends on the initial state of the volume group and whether the resource group is online or offline when the import occurs.

In all cases, the volume group will end up varied on after the resource group is started or after the cluster resources are synchronized, even if it is varied off at some point during the import process.

A failure occurs when you run the **chdev** command and have an active volume group that is in a varied on state and has the disk fencing option enabled. When you run the **chdev** command in this scenario, all SCSI3 registrations fail and the reservation for the disks in the volume group also fail. Therefore, in this scenario disk fencing is not supported for the volume group.

The following table shows the initial condition of the volume group after its creation, the state of the resource group when the import occurs, and the resulting state of the imported volume group:

Table 12. Volume group states

Initial Volume Group State	Resource Group State	Auto Imported Volume Group State
Varied ON	OFFLINE	Remains varied ON
Varied ON	ONLINE	Varied ON
Varied OFF	OFFLINE	Varied OFF until the resource group is started
Varied OFF	ONLINE	Varied ON

Importing a volume group manually:

If you do not want PowerHA SystemMirror to import your volume group automatically upon adding it to the resource group, in SMIT make sure that the PowerHA SystemMirror **Automatically Import Volume Groups** flag is set to **False** (this is the default).

To set values for the volume group:

1. Enter `smit importvg .`
2. Enter field values as follows.

Table 13. Manual import for a volume group fields

Field	Value
VOLUME GROUP name	Enter the name of the volume group that you are importing. Make sure that the volume group name is the same name that you used on the source node.
PHYSICAL VOLUME name	Enter the name of a physical volume that resides in the volume group. Note that a disk might have a different physical name on different nodes. Make sure that you use the disk name as it is defined on the destination node.
Volume Group MAJOR NUMBER	If you are using NFS, use the same major number on all nodes. Use the <code>lvlstmajor</code> command on each node to determine a free major number common to all nodes.

Changing a volume group's startup status

By default, a volume group that has just been imported is configured to automatically become active at system restart. In a PowerHA SystemMirror environment, a volume group is varied on as appropriate by the cluster event scripts. Therefore, after importing a volume group, use the Change a Volume Group panel to reconfigure the volume group so that it is not activated automatically when they system is restarted.

To change the volume group's startup status:

1. Enter `smit chvg.`
2. Select the volume group from the list.
3. Enter these field values as follows.

Table 14. Changing a volume group fields

Field	Value
Activate volume group Automatically at system restart	Set this field to No .
A QUORUM of disks required to keep the volume group online	This field is site-dependent.

Related information:

Planning shared disk and tape devices

Varying off volume group on destination nodes

Use the `varyoffvg` command to deactivate the shared volume group so that it can be imported onto another destination node or activated as appropriate by the cluster event scripts.

Enter `varyoffvg volume_group_name.`

Defining LVM components for concurrent access

These topics describe the procedure for defining shared LVM components for a concurrent access environment. Concurrent access is supported on all devices supported for PowerHA SystemMirror.

Using concurrent access with PowerHA SystemMirror requires installing an additional file set. Concurrent access does not support file systems. Instead, use logical volumes. For more information, see *Installing PowerHA SystemMirror on server nodes*.

With PowerHA SystemMirror you can create enhanced concurrent volume groups. Enhanced concurrent volume groups can be used for both concurrent and nonconcurrent access.

Take the following steps on the source and destination nodes in a PowerHA SystemMirror cluster to create a volume group that PowerHA SystemMirror can vary on in concurrent access mode.

Related concepts:

“Installing PowerHA SystemMirror on server nodes” on page 23

These topics list the prerequisites for the PowerHA SystemMirror software and describe how to install it.

Step 1. Complete prerequisite tasks

The physical volumes (hdisks) should be installed, configured, and available. You can review the disk status using the `lsdev -Cc disk` command.

Step 2. Create a concurrent access volume group on source node

The AIX operating system creates concurrent volume groups in enhanced concurrent mode by default.

To use a concurrent access volume group, create it as a *concurrent capable* volume group. A concurrent capable volume group can be activated (varied on) in either nonconcurrent mode or concurrent access mode. To define logical volumes on a concurrent-capable volume group, it must be varied on in nonconcurrent mode.

To create a concurrent-capable volume group from the command line, use the `mkvg` command. The following example creates an enhanced concurrent volume group on `hdisk1` and `hdisk2`:

```
mkvg -n -s 4 -C -y myvg hdisk1 hdisk2
```

The `mkvg` flags have the following functions.

Table 15. *Mkvg command flags*

Flag	Description
-n	Ensures that the volume group does not vary on at boot time.
-s	Specifies the partition size.
-C	Creates an enhanced concurrent volume group.
-y	Specifies the volume group name.

You can also use SMIT to run the `mkvg` command.

To create a concurrent access volume group:

1. Enter `smit cl_convg`.
2. Select **Create a Concurrent Volume Group**.
SMIT provides picklists for the node name and physical volume IDs (PVIDs) of the disks to be used. The **Add a Volume Group** panel is displayed. Some fields will already be filled and cannot be changed.
3. Enter field values as follows.

Table 16. Add a Volume Group fields

Field	Value
VOLUME GROUP name	Specify name of volume group.
Physical partition SIZE in megabytes	Accept the default.
Volume Group Major Number	Accept the default.

4. Press Enter.
SMIT prompts you to confirm your selections.
5. Press Enter.

Step 3. Import concurrent-capable volume-group information on destination nodes

Importing the volume group into the destination nodes synchronizes the definition in the PowerHA SystemMirror configuration database for the volume group on each node on which it is imported.

When you add a volume group to the resource group, you might choose to import a volume group manually onto the destination nodes or you might automatically import it onto all the destination nodes in the resource group.

Run the discovery process to make picklists available. For more information, see *Collecting information on current volume group configuration*.

The list of volume groups will include only the concurrent-capable volume groups. This list will not include **rootvg**, nor any volume groups already defined to another resource group.

Final State of Automatically Imported Volume Groups

When PowerHA SystemMirror automatically imports volume groups, the final state (varied on or varied off) depends on the initial state of the volume group and whether the resource group is online or offline when the import occurs.

In all cases, the volume group will end up varied ON after the resource group is started or after the cluster resources are synchronized, even if it is varied off at some point during the import process.

This table shows the initial condition of the volume group after its creation, the state of the resource group when the import operation occurs, and the resulting state of the imported volume group.

Table 17. Volume group and resource group states

Initial Volume Group State	Resource Group State	Auto Imported Volume Group State
Varied ON	OFFLINE	Remains varied ON
Varied ON	ONLINE	Varied ON
Varied OFF	OFFLINE	Varied OFF until the resource group is started
Varied OFF	ONLINE	Varied ON

In SMIT, if you do not want PowerHA SystemMirror to import your volume group automatically upon adding it to the resource group, make sure that the **Automatically Import Volume Groups** flag is set to **False** (this is the default), and use the **importvg** fast path.

On each destination node, manually import the volume group, using the **importvg** command, as in the following example:

```
importvg -C -y vg_name physical_volume_name
```

Specify the name of any disk in the volume group as an argument to the **importvg** command. By default, AIX automatically varies on nonconcurrent-capable volume groups when they are imported. AIX does not automatically vary on concurrent-capable volume groups when they are imported.

You can also build the **importvg** command through SMIT using the procedure that follows.

To import a concurrent-capable volume group using SMIT:

1. Enter `smit importvg`.
2. Enter field values as follows. For other fields, use the defaults or the appropriate entries for your operation.

Table 18. Import volume group fields

Field	Value
VOLUME GROUP name	Enter the name of the volume group that you are importing. Make sure the volume group name is the same name that you used on the source node.
PHYSICAL VOLUME name	Enter the name of one of the physical volumes that resides in the volume group. Note: A disk might have a different hdisk number on different nodes. Make sure that you use the disk name as it is defined on the destination node. Use the lspv command to map the hdisk number to the PVID. The PVID uniquely identifies a disk.
ACTIVATE volume group after it is imported	Set the field to No .
Volume Group MAJOR NUMBER	Accept the default.
Make this VG concurrent capable	Accept the default.
Make default varyon of VG concurrent	Accept the default.

3. Press Enter to commit the information.

If your cluster uses SCSI external disks (including RAID devices) and the import of the volume group fails, check that no reserve exists on any disk in the volume group by running the following command, only after installing the PowerHA SystemMirror software as described in Installing PowerHA SystemMirror on server nodes:

```
/usr/es/sbin/cluster/events/utls/cl_scdiskreset /dev/hdiskn ...
```

For example, if the volume group consists of *hdisk1* and *hdisk2*, enter:

```
/usr/es/sbin/cluster/events/utls/cl_scdiskreset /dev/hdisk1 /dev/hdisk2
```

Related concepts:

“Installing PowerHA SystemMirror on server nodes” on page 23

These topics list the prerequisites for the PowerHA SystemMirror software and describe how to install it.

Related tasks:

“Collecting information on current volume group configuration” on page 38

PowerHA SystemMirror can collect information about all volume groups available across all nodes in the cluster as part of the discovery process.

Step 4. Vary on the concurrent-capable volume group in nonconcurrent mode

Use the **varyonvg** command to activate a volume group in nonconcurrent mode. To create logical volumes, the volume group must be varied on in nonconcurrent access mode.

For example, to vary on the concurrent-capable volume group **myvg** in nonconcurrent access mode, enter the following command:

```
varyonvg myvg
```

To run the **varyonvg** command with SMIT, use the following procedure.

1. To vary on a concurrent capable volume group in nonconcurrent mode, enter
smit varyonvg
2. Enter field values as follows.

Table 19. Vary on volume group fields

Field	Value
VOLUME GROUP name	Specify the name of volume group.
RESYNCHRONIZE stale physical partitions	Set this field to No .
Activate volume group in SYSTEM MANAGEMENT mode	Accept the default.
FORCE activation of the volume group	Accept the default.
Varyon volume group in concurrent mode	Accept the default.

3. Press Enter.
SMIT prompts you to confirm your selections.
4. Press Enter.

Step 5. Create logical volumes on concurrent-capable volume group on source node

Create logical volumes on the volume group, specifying logical volume mirrors to provide data redundancy. If the volume group is varied on in concurrent access mode, you will not be able to create logical volumes. A concurrent-capable volume group must be varied on in nonconcurrent access mode to create logical volumes on it.

For more information about creating logical volumes, see Logical Volume Manager.

To create logical volumes on concurrent-capable volume group on source node:

1. Enter `smit cl_conlv`.
2. Specify the size of the logical volume as the number of logical partitions. Accept default values for all other fields except the following.

Table 20. Concurrent-capable volume group fields

Field	Value
Logical volume name	Specify name of logical volume. The name must be the same on all cluster nodes.
PHYSICAL VOLUME names	Specify the physical volumes you want the logical volume to include.
Number of COPIES of each logical partition	Specify 1, 2, or 3 mirror copies. Specifying 1 means no mirroring.
Mirror Write Consistency	Set this value to No .
ENABLE BAD BLOCK relocation	Set this field to No to disable bad block relocation for concurrent environments (applies to RAID devices).

Related information:

Logical Volume Manager

Step 6. Vary off volume group on source node

After creating the logical volume, vary off the volume group by using the `varyoffvg` command so that it can be varied on by the PowerHA SystemMirror scripts.

Enter `varyoffvg volume_group_name`.

Step 7. Change volume group to remain dormant at startup on destination nodes

By default, AIX configures an imported volume group to become active automatically when the system restart. In the PowerHA SystemMirror system, a volume group is varied on as appropriate by the

PowerHA SystemMirror scripts. Therefore, if you did not use the **smit cl_** fastpaths to do your configuration, you must reconfigure the volume group so that it remains dormant at startup.

To change the startup state of a volume group from the command line, enter:

```
chvg -a n volume_group_name
```

Note: The volume group must be varied on to run the **chvg** command against the volume group.

To use SMIT to change the startup state of a volume group:

1. Enter `smit chvg`.
2. Set the **Activate volume group Automatically at system restart** field to **No**. For other fields, use the defaults or the appropriate entries for your operation.
3. Press Enter to commit this change.
4. Exit SMIT.

Step 8. Complete follow-up tasks

Complete follow-up tasks as needed.

For example, ensure that the PowerHA SystemMirror scripts can activate the concurrent-capable volume group as a concurrent cluster resource.

Configuring AIX for PowerHA SystemMirror

These topics discuss several general tasks necessary for ensuring that your PowerHA SystemMirror environment works as planned.

Consult the PowerHA SystemMirror Administration guide and the PowerHA SystemMirror Performance management guide for more information on these topics.

Note: You can change the settings for a list of tunable values that were changed during the cluster maintenance and reset them to their default settings, that is, to the values that appear in the cluster after manually installing PowerHA SystemMirror.

Resetting the tunable values does not change any other aspects of the configuration, whereas installing PowerHA SystemMirror removes all user-configured configuration information including nodes, networks and resources.

Related information:

Performance management

Administering PowerHA SystemMirror

Networking considerations

These topics discuss networking-related considerations when setting up PowerHA SystemMirror.

Checking user and group IDs

If a node fails, users can still log on to the surviving nodes without experiencing problems caused by mismatches in the user or group IDs.

To avoid mismatches, make sure that user and group information is propagated to nodes as necessary. User and group IDs must be the same on all nodes.

Changing routerevalidate network option

Changing hardware and IP addresses within PowerHA SystemMirror changes and deletes routes.

Because the AIX operating system caches routes, setting the **routerevalidate** network option is required as follows:

```
routerevalidate=1
```

Updating the `/etc/hosts` file and name server configuration

Make sure all nodes can resolve all cluster addresses.

Edit the `/etc/hosts` file (and the `/etc/resolv.conf` file, if using the name server configuration) on each node in the cluster to make sure the IP addresses of all clustered interfaces are listed.

For each PowerHA SystemMirror service and boot IP label, make an entry similar to the following:

```
100.100.50.200 payroll-service
```

Also, make sure that the `/etc/hosts` file on each node has the following entry:

```
127.0.0.1 loopback localhost
```

Related information:

Using PowerHA SystemMirror with NIS and DNS

Enabling the AIX automounter daemon

For installations that require the AIX automounter daemon on PowerHA SystemMirror nodes, a modification is needed to ensure that automounter starts properly (with NFS available and running). This is due to the way PowerHA SystemMirror manages the `inittab` file and run levels upon startup.

To enable the automounter on nodes that have PowerHA SystemMirror installed, add the following line as the last line of the `/usr/es/sbin/cluster/etc/harc.net` file:

```
chnfs -g on -x 1  
startsrc -s nfsd
```

Note: The `chnfs` command change is needed only if NFSv4 is being used.

Planning PowerHA SystemMirror file collections

Certain AIX and PowerHA SystemMirror configuration files located on each cluster node must be kept in sync (be identical) for PowerHA SystemMirror to behave correctly. Such files include event scripts, application scripts, and some system and node configuration files.

Using the PowerHA SystemMirror file collection function, you can request that a list of files be kept in sync across the cluster automatically. You no longer have to copy an updated file manually to every cluster node, confirm that the file is properly copied, and confirm that each node has the same version of it. With PowerHA SystemMirror file collections enabled, PowerHA SystemMirror can detect and warn you if one or more files in a collection is deleted or has a zero value on one or more cluster nodes.

When you install PowerHA SystemMirror, it sets up the following file collections:

- **Configuration_Files**
- **HACMP_Files**

Contents of the PowerHA SystemMirror Configuration_Files collection

Configuration_Files is a container for essential system files.

These files include:

- `/etc/hosts`
- `/etc/services`
- `/etc/snmpd.conf`
- `/etc/snmpdv3.conf`

- `/etc/rc.net`
- `/etc/inetd.conf`
- `/usr/es/sbin/cluster/etc/clhosts`
- `/usr/es/sbin/cluster/etc/clinfo.rc`

PowerHA SystemMirror_Files file collection

HACMP_Files is a container for user-configuration files in the PowerHA SystemMirror configuration. PowerHA SystemMirror uses this file collection to reference all of the user-configuration files in the PowerHA SystemMirror Configuration Database classes.

The **HACMP_Files** collection automatically includes:

- Any pre-event, post-event, or notification events you have used to customize cluster events
- The start and stop scripts specified for any application controllers
- Scripts specified for application monitoring including any monitoring, notification, cleanup, and restart scripts
- Custom pager text messages
- Scripts for Tape support
- Any custom snapshot methods
- User-defined event recovery programs

Note: Do not modify or rename the PowerHA SystemMirror event script files. Also, do not include PowerHA SystemMirror event scripts in any PowerHA SystemMirror file collection.

When copying a file to a remote node, the local node's owner, group, modification time stamp, and permission settings are maintained on the remote node. That is, the remote node inherits these settings from the local node.

When you are using the default file collection remember:

- Neither of these file collections is enabled by default. If you prefer to include some user-configuration files in another file collection instead of propagating all of them, leave the **HACMP_Files** file collection disabled.
- When copying a file to a remote node, the local node's owner, group, modification time stamp, and permission settings are maintained on the remote node. That is, the remote node inherits these settings from the local node.

Permissions for all files in the **HACMP_Files** file collection is set to *execute*, which helps to prevent problems if you have not yet set execute permission for scripts on all nodes. (This is often the cause of an event failure.)

- You cannot rename the **HACMP_Files** file collection or delete it. You cannot add files to the collection or remove them. You can add a file that is already included in the **HACMP_Files** file collection (for example, an application start script) to another file collection. However, in any other case, a file can only be included in one file collection.
- You can add a file that is already included in the **HACMP_Files** file collection (for example, an application start script) to another file collection, without receiving the following error message, where *XXX_Files* is the name of the previously defined collection:
This file is already included in the <XXX_Files> collection).
- You can add and remove files or delete the **Configuration_Files** file collection.

Options for propagating a PowerHA SystemMirror file collection

Propagating a file collection copies the files in a file collection from the current node to the other cluster nodes.

Use one of the following methods to propagate a PowerHA SystemMirror file collection:

- Propagate the file collection at any time manually. You can propagate files in a file collection from the PowerHA SystemMirror File Collection SMIT menu on the local node (the node that has the files you want to propagate).
 - Set the option to propagate the file collection whenever cluster verification and synchronization is run. The node from which verification is run is the propagation node. (This is set to **No** by default.)
 - Set the option to propagate the file collection automatically after a change to one of the files in the file collection. PowerHA SystemMirror checks the file collection status on each node (every 10 minutes by default) and propagates any changes. (This is set to **No** by default.)
- One timer is set for all file collections. You can change the timer. The maximum is 1440 minutes (24 hours) and the minimum is 10 minutes.

You can set up and change file collections on a running cluster. However, note that if you add a node dynamically, the file collection on that node might have files that are not in sync with the files on the other cluster nodes. If the file collection on the node being added is set for automatic propagation upon cluster verification and synchronization, the files on the node just added are updated correctly. If this flag is not set, you must manually run the file collection propagation from one of the other nodes.

Related information:

Administering PowerHA SystemMirror

Types of error notification

Review error notifications by the AIX operating system and by PowerHA SystemMirror.

Error notification support for SIGDANGER signal

The AIX Virtual Memory Manager (VMM) uses a combination of physical memory and disk storage to provide a virtualized memory space that is much larger than the actual physical memory.

The storage-based component of this system is called the *paging space*. If not enough paging space is configured or available, the systems performance can be diminished and certain system functions might begin to fail. In a high availability environment, you want to avoid this situation because it might lead to false failovers or failures during recovery of critical resources.

When paging space becomes too low, the AIX operating system generates a SIGDANGER signal for processes that are running on the system, including the PowerHA SystemMirror subsystems. When the PowerHA SystemMirror cluster manager or clstrmgr subsystem receives a SIGDANGER signal it creates an entry in the AIX system error log. You can create a notification method that is called whenever this situation occurs. You can use this notification method to configure an automatic response when the paging space becomes too low. The response from the notification method can be as simple as sending an alert to the system administrator, or as complex as automatically recovering the page space.

To define an error notification to run when PowerHA SystemMirror receives a SIGDANGER signal, you need to create an entry in the errnotify odm class. Use the **odmadd** command with a file that contains the following information to update the odm class:

```
errnotify:
  en_name = "ha_sigdanger"
  en_persistenceflg = 1
  en_resource = "clstrmgrDANGER"
  en_method = "errpt -a -l $1 | mail -s 'SIGDANGER received' root"
```

Note: You can customize other aspects of the notification method. For example, the en_method field can specify a shell script or other executable program. For this notification method to work, you must specify the clstrmgrDANGER setting in the en_resource field.

You can change the notification method entry by using the **odmchange** command. You can remove the notification method entry by using the **odmdelete** command.

Related information:

odmchange Command
odmcreate Command
odmadd Command
Error notification

AIX error notification

Although the PowerHA SystemMirror software does not monitor the status of disk resources, it does provide a SMIT interface to the AIX error notification function. The AIX error notification function allows you to detect an event not specifically monitored by the PowerHA SystemMirror software. For example, a disk adapter failure, and to program a response to the event.

Permanent hardware errors on disk drives, controllers, or adapters might impact the fault resiliency of data. By monitoring these errors through error notification methods, you can assess the impact of a failure on the cluster's ability to provide high availability. A simple implementation of error notification would be to send a mail message to the system administrator to investigate the problem further. A more complex implementation could include logic to analyze the failure and to decide whether to continue processing, stop processing, or escalate the failure to a node failure and have the takeover node make the volume group resources available to clients.

Implement an error notification method for all errors that affect the disk subsystem. Doing so ensures that degraded fault resiliency does not remain undetected.

Note that, if you want PowerHA SystemMirror to react to a volume group failure on a node, you have an option to configure a customized AIX error notification method for this specific error, which would cause a **node_down** event or move the affected resource groups to another node.

You can customize resource recovery for a volume group that fails due to an LVM_SA_QUORCLOSE error. This error can occur if you use mirrored volume groups with quorum enabled. For this case, you can do one of the following:

- Let the PowerHA SystemMirror selective fallover function move the affected resource group
- Send a notification using the AIX Error Notification function
- Continue using your pre- and post-event scripts for this type of recovery

If you previously had a pre-event or post-event configured to handle these cases, assess how they are working with the selective fallover function. For more information about how PowerHA SystemMirror handles this particular error, see Error notification method used for volume group loss.

However, PowerHA SystemMirror does not react to any other type of volume group errors automatically. In all other cases, you still need to configure customized error notification methods, or use AIX automatic error notification methods to react to volume group failures.

For information about using this utility to assign error notification methods in one step to a number of selected disk devices, see PowerHA SystemMirror automatic error notification.

Related reference:

“PowerHA SystemMirror automatic error notification” on page 50

You can use the automatic error notification function to perform a variety of tasks.

“Error notification method used for volume group loss” on page 52

If quorum is lost for a volume group that belongs to a resource group on a cluster node, PowerHA SystemMirror selectively moves the affected resource group to another cluster node (unless you have customized resource recovery to select notify instead).

Related information:

Troubleshooting PowerHA SystemMirror

PowerHA SystemMirror automatic error notification

You can use the automatic error notification function to perform a variety of tasks.

Before you configure automatic error notification, you must have a valid PowerHA SystemMirror configuration.

Using a SMIT option, you can:

- Configure automatic error notification for cluster resources.
- List currently defined automatic error notification entries for the same cluster resources.
- Remove previously configured automatic error notification methods.

You can also use the automatic error notification function to view currently defined auto error notification entries in your PowerHA SystemMirror cluster configuration and to delete all automatic error notification methods.

Attention: The cluster must be offline when configuring automatic error notification. If the cluster is running, a warning is issued and SMIT fails.

If you add error notification methods, the AIX **cl_errnotify** utility runs automatically. This utility turns on error notification on all nodes in the cluster for the following devices:

- All disks in the **rootvg** volume group
- All disks in PowerHA SystemMirror volume groups, concurrent volume groups, and file systems
- All disks defined as PowerHA SystemMirror resources

To avoid single points of failure, the JFS log must be included in a PowerHA SystemMirror volume group.

Automatic error notification applies to selected hard, nonrecoverable error types such as disk and disk adapter. This utility does not support media errors, recovered errors, or temporary errors.

Note: You do not need to set up automatic error notification for the 2105 IBM Enterprise Storage System (ESS). These systems use the Subsystem Device Driver, which enables the hardware to handle failures itself and automatically switch to another path if it loses one.

For more information, search the IBM web site for TotalStorage support, storage software, or support for subsystem device driver or see the Support for System Storage® Multipath Subsystem Device Driver website.

Note: If you do set up automatic error notification it will simply log errors, and not initiate failover action, since the subsystem device driver handles this. However, if all PVIDs are not on VPATHS, the error notification fails. Messages are logged to the **cspoc.log** and **smit.log** files.

Executing automatic error notification assigns one of two error notification methods for all the error types noted:

- **cl_failover** is assigned if a disk or a network interface card is determined to be a single point of failure and its failure causes the cluster resources to fall over. In case of a failure of any of these devices, this method logs the error to **hacmp.out** and shuts down the cluster software on the node. It will stop cluster services with the Move Resource Groups option to shut down the node.
- **cl_logerror** is assigned for all other error types. In case of a failure of any of these devices, this method logs the error to **hacmp.out**.


The **cl_logerror** script is specified in the notification method instead of the **cl_failover** script for the following system resources:

- Disks that contain unmirrored logical volumes and therefore are considered single points of failure

- Disks that are part of volume groups or file systems defined as resources in nonconcurrent resource groups

Specifying `cl_logerror` in the notification method instead of the `cl_failover` script prevents unnecessary `node_down` events.

Related information:

 [IBM Support Portal](#)

Configuring automatic error notification:

Use this procedure to configure automatic error notification.

To configure automatic error notification:

1. Ensure that the cluster is not running.
2. Enter `smit sysmirror`.
3. In SMIT, select **Problem Determination Tools > PowerHA SystemMirror Error Notification > Configure Automatic Error Notification**. The SMIT menu contains the following items:

Table 21. Configure Automatic Error Notification fields

Field	Value
List Error Notify Methods for Cluster Resources	Lists all currently defined automatic error notification entries for certain cluster resources: PowerHA SystemMirror defined volume groups, concurrent volume groups, file systems, and disks; rootvg . The list is shown to the screen.
Add Error Notify Methods for Cluster Resources	Error notification methods are automatically configured on all relevant cluster nodes.
Delete Error Notify Methods for Cluster Resources	Error notification methods previously configured with the Add Error Notify Methods for Cluster Resources option are deleted on all relevant cluster nodes.

4. Select the **Add Error Notify Methods for Cluster Resources** option from the list.
5. Optional: Because error notification is automatically configured for all the listed devices on all nodes, make any modifications to individual devices or nodes manually, after running this utility.
If you make any changes to cluster topology or resource configuration, you might need to reconfigure automatic error notification. When you run **verification** after making any change to the cluster configuration, you will be reminded to reconfigure error notification if necessary.

Listing error notification methods:

You can use PowerHA SystemMirror to view automatic error notification methods that exist for your cluster configuration.

To list error notification methods, complete the following steps:

1. Enter `smit sysmirror`.
2. In SMIT, select **Problem Determination Tools > PowerHA SystemMirror Error Notification > Configure Automatic Error Notification** and press Enter.
3. Select the **List Error Notify Methods for Cluster Resources** option. The utility lists all currently defined automatic error notification entries with these PowerHA SystemMirror components: PowerHA SystemMirror-defined volume groups, concurrent volume groups, file systems, and disks; **rootvg**. The following example shows output for cluster nodes named *sioux* and *quahog*:

COMMAND STATUS

Command: OK stdout: yes stderr: no
Before command completion, additional instructions may appear below.

```

sioux:
sioux: PowerHA SystemMirror Resource      Error Notify Method
sioux:
sioux: hdisk0 /usr/es/sbin/cluster/diag/cl_failover
sioux: hdisk1 /usr/es/sbin/cluster/diag/cl_failover
sioux: scsi0  /usr/es/sbin/cluster/diag/cl_failover
quahog:
quahog: PowerHA SystemMirror Resource      Error Notify Method
quahog:
quahog: hdisk0 /usr/es/sbin/cluster/diag/cl_failover
quahog: scsi0  /usr/es/sbin/cluster/diag/cl_failover

```

Deleting error notification methods:

Use this procedure to delete automatic error notification entries previously assigned by using this function.

Take the following steps:

1. Enter `smit sysmirror`.
2. In SMIT, select **Problem Determination Tools > PowerHA SystemMirror Error Notification > Configure Automatic Error Notification** and press Enter.
3. Select the **Delete Error Notify Methods for Cluster Resources** option. Error notification methods previously configured with the **Add Error Notify Methods for Cluster Resources** option are deleted on all relevant cluster nodes.

PowerHA SystemMirror monitoring system events

In PowerHA SystemMirror some system events are automatically monitored with the new event manager daemon.

The event manager daemon automatically monitors for system events.

The following events are monitored by the event manager daemon:

- Cluster Aware AIX repository disk online or offline
- Volume Group online or offline (for volume groups that are managed by PowerHA SystemMirror)
- rootvg online or offline

You can monitor the rootvg events only when the rootvg disk is using the native AIX Multipath I/O (MPIO) driver and the rootvg disk is not an internal parallel SCSI disk. To verify whether the rootvg disk is using the MPIO driver, on the command line, type `lspath -l hdiskname`, where *hdiskname* is the name of the rootvg disk. If the rootvg disk is not using the MPIO driver, the following error message is displayed:

```

lspath: 0514-538 Cannot perform the requested function because the
        specified device does not support multiple paths.

```

To test the notification and response to rootvg loss, you must stimulate the I/O to the rootvg.

Responses to monitored events are configured by default. You can change event responses by using the SMIT menu, **Custom Cluster Configuration > Events > System Events**.

Error notification method used for volume group loss

If quorum is lost for a volume group that belongs to a resource group on a cluster node, PowerHA SystemMirror selectively moves the affected resource group to another cluster node (unless you have customized resource recovery to select notify instead).

For this action, PowerHA SystemMirror uses an automatic error notification method to inform the Cluster Manager about the failure of a volume group. The system checks whether the `LVM_SA_QUORCLOSE`

error appeared in the AIX error log file on a cluster node and informs the Cluster Manager to selectively move the affected resource group. PowerHA SystemMirror uses this error notification method only for mirrored volume groups with quorum enabled.

You do not need to set up automatic error notification for the 2105 IBM Enterprise Storage System (ESS). These systems use the Subsystem Device Driver, which enables the hardware to handle failures itself and automatically switch to another path if it loses one.

If you do set up automatic error notification, it will simply log errors and not initiate failover action since the Subsystem Device Driver handles this. However, if all PVIDs are not on VPATHS, the error notification fails. Messages are logged to the **csnoc.log** and to the **smit.log**.

Note: Do not modify the error notification method used by PowerHA SystemMirror to react to a volume group loss. PowerHA SystemMirror issues a warning and takes no action if you attempt to customize this notification method or use it to protect against the failure of other types of resources.

The automatically configured AIX Error Notification method is launched if it finds:

- The error LVM_SA_QUORCLOSE in the AIX error log on a cluster node.
- The appropriate entries in the **errnotify** class in the PowerHA SystemMirror configuration database on that node. The **errnotify** entries are created during synchronization of the cluster resources.

The AIX Error Notification method updates the Cluster Manager. The Cluster Manager then tries to move the resource group that has been affected by a volume group loss to another node in the cluster.

If failover does not occur, check that the LVM_SA_QUORCLOSE error appeared in the AIX error log. When the AIX error log buffer is full, new entries are discarded until space becomes available in the buffer, and, therefore, AIX error notification does not update the Cluster Manager to selectively move the affected resource group.

Note: You can change the default selective failover action to be a notify action instead. For more information, see the Administration guide.

If the AIX **errdaemon** is not running on a cluster node, PowerHA SystemMirror has no means to detect the loss of quorum error in the AIX log file, and therefore, cannot selectively move a resource group if it contains a failed volume group. In this case, PowerHA SystemMirror issues a warning.

The automatically configured AIX error notification method works correctly if the following requirements are met:

- Do not modify this error notification method.
- Synchronize the cluster after making changes to the cluster configuration. A notification script used for a volume group failure should correspond to the current configuration of cluster resources. Otherwise, PowerHA SystemMirror issues a warning during verification and takes no action to selectively move the affected resource group.
- Besides the **errnotify** entries created by PowerHA SystemMirror for selective failover, the **errnotify** class in the PowerHA SystemMirror configuration database might also contain other entries related to the same AIX error labels and resources. However, the selective failover utility provides the most effective recovery mechanism to protect a resource group from the failure of a single resource.

The notification method that is run in the case of a volume group failure provides the following information in the **hacmp.out** log file:

- AIX error label and ID
- The name of the affected system resource (resource group)
- The node's name on which the error occurred

Related information:

Resource group behavior during cluster events

Test error notification methods

You can test the error notification methods generated by the selective fallover facility by emulating an error for each volume group in SMIT.

To test error notification:

1. Enter `smit sysmirror`.
2. In SMIT, select **Problem Determination Tools > PowerHA SystemMirror Error Notification > Emulate Error Log Entry**.
3. Select from the picklist the error notification object that was generated by the selective fallover facility for each volume group.

Emulation of error log entries

After you have added one or more error notification methods to the AIX Error Notification facility, test your methods by emulating an error.

By inserting an error log entry into the AIX error device file (`/dev/error`), you cause the AIX error daemon to run the appropriate specified notify method. This enables you to determine whether your predefined action is carried through.

To emulate an error log entry:

1. Enter `smit sysmirror`.
2. In SMIT, select **Problem Determination Tools > PowerHA SystemMirror Error Notification > Emulate Error Log Entry**.

The **Select Error Label** shows a picklist of the notification objects for which notification methods have been defined. The list includes error notification objects generated by both the Automatic Error Notification facility and by the selective fallover facility for volume group loss. For volume group loss for the description of these methods, See error notification method used.

3. Select a notification object and press Enter to begin the emulation.

As soon as you press Enter, the emulation process begins: The emulator inserts the specified error into the AIX error log, and the AIX error daemon runs the notification method for the specified object.

When the emulation is complete, you can view the error log by typing the `errpt` command to be sure the emulation took place. The error log entry has either the resource name EMULATOR or a name the user specified in the **Resource Name** field during the process of creating an error notify object.

You will now be able to determine whether the specified notify method was carried out.

Note: Remember that the actual notify method will be run. Whatever message, action, or executable you defined will occur. Depending on what it is, you might need to take some action.

Related reference:

“Error notification method used for volume group loss” on page 52

If quorum is lost for a volume group that belongs to a resource group on a cluster node, PowerHA SystemMirror selectively moves the affected resource group to another cluster node (unless you have customized resource recovery to select notify instead).

OEM disk, volume group, and file systems in a cluster

These topics describe how you can customize PowerHA SystemMirror software to integrate original equipment manufacturer (OEM) disks, volume groups, and file systems in a PowerHA SystemMirror cluster.

Integrating OEM disks in a PowerHA SystemMirror cluster

You can use IBM or OEM storage disks with PowerHA SystemMirror to create a highly available infrastructure. The storage disks must be defined in AIX and be part of an AIX LVM volume group.

Depending on the type of OEM disk, you can use custom methods or an OEM disk vendor methods to tell PowerHA SystemMirror that an unknown disk should be treated the same way as a known and supported disk type. You can also use these methods to specify the custom methods that provide the low-level disk processing functions supported by PowerHA SystemMirror for that particular disk type.

Custom methods for OEM disks provide a means for configuring and supporting OEM disks in a PowerHA SystemMirror cluster. These methods depend on the types of disks that are supported by the PowerHA SystemMirror software, and on the different ways they are handled. A level of similarity between certain types of disks allows you to configure an OEM disk that is unknown to PowerHA SystemMirror as you would a similar type of disk known to PowerHA SystemMirror. Alternatively, you can select one or more of the disk processing methods used by PowerHA SystemMirror to safely deal with an OEM disk.

The disk processing methods include:

- Identifying ghost disks
- Ghost disks are duplicate profiles of the disks created during disk configuration processing. Ghost disks must be removed to allow AIX to vary on the volume group residing on the original disks
- Determining whether another node in the cluster is holding a disk reserve
- Breaking a disk reserve
- Making a disk available for use by another node

Handling disks in PowerHA SystemMirror

PowerHA SystemMirror handles disks according to AIX LVM disk requirements.

The AIX Logical Volume Manager (LVM) by default is not designed to handle multiple nodes accessing the same set of disks, which are referred to as shared. Therefore, PowerHA SystemMirror performs special disk processing for shared disks in a cluster. When nodes in the PowerHA SystemMirror cluster join and leave the cluster, PowerHA SystemMirror brings the shared disks into a state where the LVM commands work correctly.

In the LVM, one or more disks, also known as physical volumes, can be grouped together to form a volume group. Logical volumes can be built on top of a volume group and file systems can be built on top of logical volumes. See AIX LVM documentation for more information.

PowerHA SystemMirror supports two modes of access to volume groups in a cluster: serial access mode and concurrent access mode.

Concurrent access to volume groups:

In concurrent access, the volume groups residing on shared disks are activated and can be simultaneously accessed in read/write mode by all active nodes in the cluster.

Note: AIX introduced enhanced concurrent mode. When concurrent volume groups are created on AIX, they are automatically created as enhanced concurrent mode. Convert your RAID concurrent volume groups to enhanced concurrent mode volume groups.

Accessing disks in concurrent mode requires that software packages running above the LVM coordinate their operations across the cluster to ensure data integrity. The AIX Journaled File System is not supported on volume groups accessed in concurrent mode. Consult the provider of any middleware to determine whether that middleware can reliably function in concurrent mode.

For concurrent access volume groups, PowerHA SystemMirror must ensure that no disk reserve is placed on the disk so that access by other nodes is possible.

The question arises as to whether any disk can be used in concurrent mode. The answer is no, because of a practical requirement to mirror data in a PowerHA SystemMirror cluster, so that data is not be lost if any of the disks fails. Even though most disks could be accessed without establishing a disk reserve on them, many disks still cannot be used in concurrent mode, because the LVM cannot simultaneously update all the copies of a single logical partition. Even when the write operation on a disk are done in parallel, they might not actually be completed at the same time, due to other activity on the disk or due to the physical characteristics of the disks. If one node writes data and the other node reads it, there is no guarantee that the reader gets the latest copy. Additionally, if the LVM cannot update its structures on a disk, it ceases to write to that disk and marks it as failed. In a multi-node cluster, this situation creates a single point of failure.

Related information:

Planning shared LVM components

Disks supported by PowerHA SystemMirror

PowerHA SystemMirror supports enhances concurrent mode disks and serial access mode disks.

Enhanced concurrent mode disks

The AIX operating system supports enhanced concurrent mode. In enhanced concurrent mode, instances of the Concurrent Logical Volume Manager (CLVM) coordinate changes between nodes through the Group Services component of the Reliable Scalable Cluster Technology (RSCT) facility in AIX. Group Services protocols flow over the communications links between the cluster nodes.

With enhanced concurrent mode:

- Any disk supported by PowerHA SystemMirror for attachment to multiple nodes can be included in an enhanced concurrent mode volume group.
- Enhanced concurrent mode volume groups require the Concurrent Resource Manager. This provides the cluster initialization that is needed for CLVM to provide enhanced concurrent mode support.

Serial access mode disks

Only one node at a time can access serially accessed disks in read/write mode, and the shared volume groups residing on these disks can be varied on only one node in a cluster. The LVM perceives these disks as normal disks. The actual hdisks as perceived by the LVM can be either physical spindles (a configuration sometimes referred to as JBOD (Just a Bunch of Disks) to contrast it with RAID) or RAID disks. For JBOD configurations, configure the LVM to provide mirroring for data reliability.

Related information:

 Strategic and Innovative RAID solutions

Supporting disks with known characteristics

All the disks supported by PowerHA SystemMirror adhere to some form of the SCSI standard. However, this standard has undergone considerable evolution over time and allows for many implementation choices. Even so, many disks behave in similar ways. PowerHA SystemMirror provides mechanisms that will allow you, while configuring a cluster, to direct PowerHA SystemMirror to treat an unknown disk exactly the same way as another disk it supports.

Use the `/etc/cluster/lunreset.lst` file to tell PowerHA SystemMirror that a particular disk supports LUN resets, even though its response to the SCSI inquiry command indicates that the disk does not support the SCSI command set. The file contains the **Vendor Identification** field. A SCSI inquiry command returns the value of this field. Consult the disk manufacturer for the value of this field.

Use the `/etc/cluster/disktype.lst` file to tell PowerHA SystemMirror that it can process a particular disk the same way it processes disks that it supports. The file contains a series of lines of the following form:
`<PdDvLn field of the hdisk><tab><supported disk type>`

To determine the value of the `PdDvLn` field for a particular `hdisk`, enter the following command:

```
lsdev -Cc disk -l <hdisk name> -F PdDvLn
```

This command returns a response similar to the following command:

```
disk/scsi/scsd
```

The supported disk types are.

Table 22. Supported disk types

Disk name in PowerHA SystemMirror	Disk type
SCSIDISK	SCSI-2 Disk
FCPARRAY	Fibre Attached Disk Array
ARRAY	SCSI Disk Array
FSCSI	Fibre Attached SCSI Disk

PowerHA SystemMirror does not modify the previously described files after they have been configured and are not removed if the product is uninstalled. This ensures that customized modifications are unaffected by the changes in PowerHA SystemMirror. By default, the files initially contain comments explaining their format and usage.

Keep in mind that the entries in these files are classified by disk type, not by the number of disks of the same type. If there are several disks of the same type attached to a cluster, there should be only one file entry for that disk type. In addition, unlike other configuration information, PowerHA SystemMirror does not automatically propagate these files across nodes in a cluster. It is your responsibility to ensure that these files contain the appropriate content on all cluster nodes. Use the PowerHA SystemMirror File Collections facility to propagate this information to all cluster nodes.

Related information:

Managing PowerHA SystemMirror file collections

Customizing PowerHA SystemMirror disk processing

Some disks might behave sufficiently differently from those supported by PowerHA SystemMirror so that it is not possible to achieve correct results by telling PowerHA SystemMirror to process these disks exactly the same way as supported disk types. For these cases, PowerHA SystemMirror provides finer control. While doing cluster configuration, you can either select or specify one of the specific methods to be used for the steps in disk processing.

PowerHA SystemMirror supports the following disk processing methods:

- Identifying ghost disks
- Determining whether a disk reserve is being held by another node in the cluster
- Breaking a disk reserve
- Making a disk available for use by another node

Using PowerHA SystemMirror you can specify any of its own methods for each step in disk processing or to use a customized method, which you define. This allows you to use a custom-defined method only for configuring a disk, and to use the existing PowerHA SystemMirror methods for other disk operations.

Determining whether a disk reserve is held by another node:

A disk reserve restricts access to the disk by a specific node. If the node that established the reserve has failed, PowerHA SystemMirror takes special steps to remove that disk reserve.

This method is passed as an input *hdisk* name, such as *hdisk42*. It passes back a return code with one of the following values:

- 0 The disk is not reserved, and it is readable and writable.
- 1 A disk reserve is currently held by another node.
- 2 An error has occurred. PowerHA SystemMirror should terminate processing for this disk.

Any other value will be interpreted as equivalent to 2.

PowerHA SystemMirror supports a value of *SCSI_TUR* for the name of this method. If this value is specified, existing processing in PowerHA SystemMirror is used to determine whether a disk reserve is being held. The processing opens the disk device and sends a SCSI Test Unit Ready command via `ioctl()`. A response of reservation conflict means that there is a disk reserve held by another node.

Breaking a disk reserve:

Disk reserves can be broken in parallel. Any serialization requirements must be handled by user-supplied methods.

The following method is passed as input for two formal parameters:

- The name of the owning adapter for the disk, such as *scsi1*. This is determined by PowerHA SystemMirror by running the command:
`lsdev -Cc disk -l <hdisk name> -F parent`
- The *hdisk* name, such as *hdisk42*.

A return code of 0 indicates success; any other value causes PowerHA SystemMirror to terminate processing for this disk.

PowerHA SystemMirror supports values of *TARGET*, *PSCSI*, and *FSCSI* for the name of this method. If one of these values is specified, PowerHA SystemMirror uses existing processing for this method. The processing breaks the disk reserves for these types of disks:

- *TARGET*. A SCSI target ID reset will be sent by executing
`openx(<hdisk name>, ,SC_FORCED_OPEN)`
Specifying the *TARGET* value is appropriate for SCSI devices that have only one LUN per SCSI ID.
- *PSCSI*. The disk is treated as a parallel SCSI device. In particular, the SCSI and LUN ID information is retrieved from the position attribute in the **CuDV** entry for the *hdisk*. The disk is opened, and either a LUN reset or a target reset is sent, depending on whether a SCSI inquiry command reports this device as being *SCSI-2* or *SCSI-3*. Note, that an entry in `/etc/cluster/lunreset.lst` causes a LUN reset to be sent to a device that identifies itself as *SCSI-2*.
- *FSCSI*. The disk is treated as a fiber-attached SCSI device. In particular, the LUN ID is retrieved from the `lun_id` field in the **CuAt** entry for this device.

Making a disk available for use by another node:

This method allows for making a disk available in the AIX state, as returned by `lsdev` command. A device must be in an available state for the LVM to be able to access it and to vary on a volume group.

This method is passed as an input *hdisk* name, such as *hdisk42*.

A return code of zero indicates success; any other value causes PowerHA SystemMirror to terminate processing for this disk.

PowerHA SystemMirror supports a value of *MKDEV* for the name of this method. If this value is specified, PowerHA SystemMirror uses existing processing for this method and attempts to run the following command up to four times:

```
mkdev -l <hdisk name>
```

Configuring OEM disk methods in SMIT

You can add, change, and remove custom disk processing methods for a specific OEM disk using SMIT. You can select existing PowerHA SystemMirror-supported custom disk methods or use your own custom disk methods.

Note: There are no SMIT panels for modifying the following files:

/etc/cluster/conraid.dat

/etc/cluster/lunreset.lst

/etc/cluster/disktype.lst

These are text files that can be modified with an editor.

Related reference:

“Supporting disks with known characteristics” on page 56

All the disks supported by PowerHA SystemMirror adhere to some form of the SCSI standard. However, this standard has undergone considerable evolution over time and allows for many implementation choices. Even so, many disks behave in similar ways. PowerHA SystemMirror provides mechanisms that will allow you, while configuring a cluster, to direct PowerHA SystemMirror to treat an unknown disk exactly the same way as another disk it supports.

Adding custom disk methods:

You can add custom disk methods using SMIT.

To add a custom disk method:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Resources > Custom Disk Methods > Add Custom Disk Methods**, and press Enter.
3. Enter field values that define the disk processing methods you want to specify for the particular OEM disk:

Table 23. Add Custom Disk Methods fields

Field	Value
Disk type (PdDvLn field from CuDv)	Enter the identifier for the particular disk type. It is the value of the PdDvLn field of the CuDv entry for the disk.
Method to identify ghost disks	<p>You can select a method from the picklist, or you can enter the full path name of a custom method that PowerHA SystemMirror can use to identify ghost disks. Select the appropriate method, and press Enter.</p> <p>Ghost disks are duplicate profiles of the disks created during disk configuration processing. Ghost disks must be removed to allow the AIX operating system to vary on the volume group that resides on the original disks.</p>

Table 23. Add Custom Disk Methods fields (continued)

Field	Value
Method to determine if a reserve is held	<p>You can select a method from the picklist, or you can enter the full path name of a custom method that PowerHA SystemMirror can use to determine whether another node holds a reserve on the specified disk. Select the appropriate method, and press Enter.</p> <p>A disk reserve restricts access to the disk by a specific node. If the node that established the reserve fails, PowerHA SystemMirror removes that disk reserve.</p>
Method to break a reserve	<p>You can select a method from the picklist, or you can enter the full path name of a custom method that PowerHA SystemMirror can use to break a reserve on the specified disk. Select the appropriate method, and press Enter.</p> <p>A disk reserve restricts access to the disk by a specific node. If the node that established the reserve fails, PowerHA SystemMirror removes that disk reserve.</p>
Break reserves in parallel	<p>Select true or false. The default is false.</p> <p>Some disk processing methods can be safely run in parallel. This method might provide a performance advantage in cluster configurations with a large number of disks.</p>
Methods to make the disk available	<p>You can select a method from the picklist, or you can enter the full path name of a custom method that PowerHA SystemMirror can use to make this disk available. Select the appropriate method, and press Enter.</p> <p>After a disk becomes accessible and any reserve is removed from that disk, it must be made available for the AIX operating system to access that disk.</p>

Note: The custom disk processing method that you specify for a particular OEM disk is added only to the local node. This information is not propagated to other nodes. You must copy this custom disk processing method to each node manually or use the PowerHA SystemMirror File Collections facility. After you make the selections, the information is applied to the disks defined on the local node.

4. Configure the same custom disk processing method on each other node in the cluster and synchronize the cluster resources. The cluster verification process ensures that the method that you configured exists and is executable on all nodes. The synchronization process ensures that the entries in the PowerHA SystemMirror configuration database are the same on all nodes, but this process does not synchronize the methods named in the database entries.

Changing custom disk methods:

You can change custom disk methods using SMIT.

To change a custom disk method:

1. Enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Resources > Custom Disk Methods > Change/Show Custom Disk Methods** and press Enter.
SMIT displays a picklist with the names of the specified disk processing methods.
3. Select a name of a particular disk method and press Enter. SMIT displays the current information.
4. Enter new information in the fields you want to change and press Enter.

Note: The characteristics of the method are changed on the local node but they are not updated on remote nodes. Custom disk methods should be updated manually on other nodes in the cluster, or you can use the PowerHA SystemMirror File Collections facility for this purpose.

Removing custom disk methods:

You can use SMIT to remove custom disk methods.

To remove a custom disk method, complete the following steps:

1. Enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Resources > Custom Disk Methods > Remove Custom Disk Methods**, and press Enter.
SMIT displays a picklist with the names of the specified disk processing methods.
3. Select a disk method that you want to remove, and press Enter. SMIT prompts you to confirm your selection.

Note: PowerHA SystemMirror deletes the characteristics of the disk methods on the local node but does not update them on remote nodes. You have to manually delete the custom disk methods on the remote nodes in the cluster.

Integrating OEM volume groups in a PowerHA SystemMirror cluster

You can configure OEM volume groups in AIX and use PowerHA SystemMirror as an IBM High Availability solution to manage such volume groups, their corresponding file systems, and application controllers.

Application controllers are defined as start and stop scripts for the applications supported by the volume groups.

For information on file systems, see *Integrating OEM file systems in a PowerHA SystemMirror cluster*.

Note: Different OEMs may use different terminology to refer to similar constructs. From now on, this documentation uses the term *volume groups* to refer to OEM and Veritas volume groups.

In particular, PowerHA SystemMirror automatically detects and provides the methods for volume groups created with the Veritas Volume Manager (VxVM) using Veritas Foundation Suite (VFS) v. 4.0. Note, Veritas Foundation Suite is also referred to as Veritas Storage Foundation (VSF). This documentation uses VFS.

Use the following table to identify the terminology used for the storage components by IBM and Veritas. Other manufacturers might use similar terminology:

Table 24. Storage components terminology

AIX Logical Volume Manager (LVM)	Veritas Volume Manager (VxVM)
Physical Volumes	Disk Media
Logical Volumes	Volumes
Volume Groups	Disk Groups

Related concepts:

“Integrating OEM file systems in a PowerHA SystemMirror cluster” on page 66

You can configure OEM file systems in the AIX operation system and use PowerHA SystemMirror as an IBM High Availability solution to manage such volume groups, their corresponding file systems, and application controllers.

Overview for integrating OEM volume group functions

Depending on the type of OEM volume, custom methods in PowerHA SystemMirror allow you (or an OEM vendor) either to tell PowerHA SystemMirror that a volume unknown to AIX LVM should be treated the same way as a known and supported volume or to specify the custom methods that provide the volume groups processing functions supported by PowerHA SystemMirror.

You can either use custom methods for volume groups offered by PowerHA SystemMirror, or create your own custom methods to use for the non-IBM volume groups in the cluster. These functions are performed within the normal PowerHA SystemMirror event processing.

Prerequisites and limitations:

In order for PowerHA SystemMirror to handle OEM volume groups and file systems in your configuration, the volume groups and file systems must include these prerequisites. There also are some limitations.

Volume groups and file systems must adhere to these conditions:

- Each OEM volume is composed of one or more physical disks (equivalent to `hdisks` in the AIX operating system) and the system can determine the names of the physical disks from the name of the volume.
- OEM disks, volume groups, and file systems must have operations and sequences of operations comparable to the functions of AIX LVM, although the command names, syntax, and arguments may differ.
- For Veritas volume groups, PowerHA SystemMirror automatically uses them once they are configured in PowerHA SystemMirror with the SMIT panels described below. For other OEM volume groups, PowerHA SystemMirror uses the custom methods that you specify in SMIT.

Limitations

PowerHA SystemMirror identifies OEM volume groups and file systems and performs all major functions with them. See OEM Volume group functions performed by PowerHA SystemMirror and OEM file systems functions performed by PowerHA SystemMirror.

However, some of the PowerHA SystemMirror functions have limitations for OEM volume groups and file systems:

- You cannot use C-SPOC (Cluster Single Point of Control) to manage OEM disk, volume, and file systems operations. In particular, if you use C-SPOC for your other operations, PowerHA SystemMirror does not include OEM disks, volume groups and file systems in picklists for your selection in SMIT.
- PowerHA SystemMirror does not automatically discover OEM volume groups and file systems and does not list them for your selections in picklists in SMIT.
- PowerHA SystemMirror does not use NFS (Network File System) functions for OEM file systems.
- PowerHA SystemMirror does not provide a workaround for any limitations of OEM disks, volume groups, and file systems that exist outside of PowerHA SystemMirror.
- In addition to listing, varying on and off, and verifying volume groups and file systems created with the AIX LVM, PowerHA SystemMirror supports other extended functions for its "own" volume groups and file systems, such as enhanced concurrent mode, active and passive varyon process, heartbeating over disk, selective failover upon volume group loss and other.

These functions of PowerHA SystemMirror utilize the AIX LVM capabilities, and since OEM volume groups and file systems do not use AIX LVM, PowerHA SystemMirror does *not* support these functions for OEM volume groups.

- PowerHA SystemMirror mechanism for fast disk takeover cannot be utilized for mixed physical volumes, that is, for disks comprised of both IBM and non-IBM disks. The LVM enhanced concurrent mode capabilities of the AIX operating system are used to enable fast disk takeover. Only LVM supported disks can be used with enhanced concurrent mode and the fast disk takeover feature.
- The automatic error notification methods available in the AIX operating system cannot be configured in PowerHA SystemMirror for OEM volume groups and file systems.
- You can define sites in the cluster, but you cannot configure and use replicated resources for PowerHA SystemMirror Enterprise Edition in the cluster that uses OEM volume groups. The intersite management policy for resource groups must be set to Ignore.
- PowerHA SystemMirror does not support the LVM cross-site mirroring function for resource groups that have OEM volume groups and file systems.

Software requirements

PowerHA SystemMirror supports Veritas volume groups at whatever level of VFS software is running on the system (depending on the version level supported by VFS 4.0).

Related concepts:

“OEM file systems functions performed by PowerHA SystemMirror” on page 68

When PowerHA SystemMirror identifies OEM file systems of a particular type, it provides the processing functions for them or you can specify your own custom methods for any one of these functions.

Related reference:

“OEM Volume group functions performed by PowerHA SystemMirror”

When PowerHA SystemMirror identifies OEM volume groups of a particular type, it provides built-in processing functions for them or allows you to specify your own custom methods for any one of these functions.

OEM Volume group functions performed by PowerHA SystemMirror:

When PowerHA SystemMirror identifies OEM volume groups of a particular type, it provides built-in processing functions for them or allows you to specify your own custom methods for any one of these functions.

These functions include:

- List volume groups of a specified type defined on the nodes.
- List physical and logical disks comprising a specified volume group.
- Bring a volume group online and offline.
- Determine a volume group status.
- Verify volume groups configuration.
- Provide a location of log files and other debugging information. You can then view this information by using the AIX **snap -e** command.

PowerHA SystemMirror enables you to manage OEM volume groups in a PowerHA SystemMirror cluster by bringing them online and offline as needed for cluster event processing.

Before bringing a volume group online or offline, PowerHA SystemMirror checks the volume group type to determine whether it must use the corresponding built-in custom method to manage the volume group and file system and bring it online or offline as needed.

OEM volume groups and file systems as resources in a PowerHA SystemMirror resource group:

You can include OEM disks, volume groups, and file systems in a PowerHA SystemMirror resource group. PowerHA SystemMirror recognizes OEM volume groups and file systems and handles them as the AIX LVM volume groups and file systems by listing, activating, checking and taking them offline when it is necessary for the cluster event processing tasks.

You can also mix the OEM disks, volume groups and file systems with AIX volume groups and file systems in a single PowerHA SystemMirror resource group.

When you include OEM or Veritas volume groups or file systems into PowerHA SystemMirror resource groups, **Automatically Import Volume Groups** field should be set to **False**. Also OEM volume groups must be set to not automatically varyon when the node is restarted. OEM file systems must be set to not automatically mount when the node is restarted.

To view OEM disks, volume groups and file systems used in the PowerHA SystemMirror cluster, you can use the PowerHA SystemMirror SMIT interface.

Supporting Veritas volume groups

Among other OEM volume groups and file systems, PowerHA SystemMirror supports volume groups and file systems created with VxVM in Veritas Foundation Suite v.4.0.

To make it easier for you to accommodate Veritas volume groups in the PowerHA SystemMirror cluster, the methods for Veritas volume groups support are predefined in PowerHA SystemMirror and are used automatically. After you add Veritas volume groups to PowerHA SystemMirror resource groups, you can select the methods for the volume groups from the picklists in PowerHA SystemMirror SMIT menus for OEM volume groups support.

Note: After you configure Veritas volume groups or file systems in PowerHA SystemMirror cluster, the PowerHA SystemMirror cluster verification utility runs the associated Veritas verification check. This verification check is effective only for the volume groups or file systems that are currently online on the node on which the verification is being run.

Configuring OEM volume groups methods in SMIT

You can add, change, and remove custom volume groups processing methods for a specific OEM volume group using SMIT. You can select existing custom volume group methods that are supported by PowerHA SystemMirror, or you can use your own custom methods.

You can add, change, and remove the methods that PowerHA SystemMirror will use to manage OEM volume groups.

Adding custom volume group methods:

You can add custom volume group methods using SMIT.

To add a custom volume group method:

1. Enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Resources > Custom Volume Group Methods > Add Custom Volume Group Methods**, and press Enter.

Note: If you are configuring Veritas volume groups, perform this operation in SMIT on the same node on which the volume group is currently active.

3. Enter field values that define the volume group processing methods you want to specify for the particular OEM volume group:

Table 25. Add Custom Volume Group Methods fields

Field	Value
Volume Group Type	Enter the identifier for the particular volume group type, or use your own custom volume group methods. For Veritas volume groups, PowerHA SystemMirror uses the method to call this type. By default, it is the value of the PdDvLn field of the CuDv entry for the volume.
Method to List Volume Group Names	Select a method from the picklist, or enter the full path name of a custom method that PowerHA SystemMirror should use to list volume group names for the specified volume group. As a default method, PowerHA SystemMirror uses the AIX lsvg command. The custom method that you use must generate a list of the volume group names as an output. PowerHA SystemMirror expects that the method returns zero for success and non-zero for failure. PowerHA SystemMirror also expects that if a method passed to it is a list, the list should have one item per line.

Table 25. Add Custom Volume Group Methods fields (continued)

Field	Value
Method to determine physical disks (hdisks) comprising the volume group	<p>Select a method from the picklist, or enter the full path name of a custom method that PowerHA SystemMirror should use to determine which physical disks must be made available on the node before PowerHA SystemMirror brings the specified volume group online.</p> <p>As a default method, PowerHA SystemMirror uses LSPV which calls the AIX lspv command.</p> <p>The custom method that you use must generate a list of the <code>hdisk</code> names as an output. PowerHA SystemMirror expects that the method returns zero for success and non-zero for failure.</p>
Method for bringing the volume group online	<p>Select a method from the picklist, or enter the full path name of a custom method that PowerHA SystemMirror should use to bring the volume group online.</p> <p>As a default method, PowerHA SystemMirror uses VARYONVG which calls the AIX varyonvg command.</p> <p>PowerHA SystemMirror expects that the method returns zero for success and non-zero for failure.</p>
Method for determining volume group status	<p>Select the method from the picklist, or enter the full path name of a custom method that PowerHA SystemMirror should use to determine the volume group status.</p> <p>As a default method, PowerHA SystemMirror uses LSACTIVEVG which calls the AIX lsvg -o command.</p> <p>PowerHA SystemMirror takes as an input the volume group name and expects that the method returns 0 for offline, 1 for online and 2 for command failure. PowerHA SystemMirror also expects that if a method passed to it is a list, the list should have one item per line.</p>
Method for bringing the volume group offline	<p>Select a method from the picklist, or enter the full path name of a custom method that PowerHA SystemMirror should use to take the volume group offline.</p> <p>As a default method, PowerHA SystemMirror uses VARYOFFVG which calls the AIX varyoffvg command.</p> <p>PowerHA SystemMirror expects that the method returns zero for success and non-zero for failure.</p>
Method for verifying volume configuration	<p>Enter the full path name of a custom method that PowerHA SystemMirror should use to verify the volume configuration on each cluster node.</p> <p>Note: By default, PowerHA SystemMirror verifies the configuration for AIX volume groups and file systems. It also uses the Veritas verification method for Veritas volume groups and file systems. However, the Veritas verification is effective <i>only</i> for those volume groups and file systems that are currently active on the nodes on which the verification is being run. PowerHA SystemMirror expects that the method returns zero for success and non-zero for failure.</p>
Directory containing log information (optional)	<p>Empty by default. Enter one or more space-separated directories in this field. You can use the AIX snap -e command to view this information.</p>

Note: The custom volume group processing method that you specify for a particular OEM volume group is added to the local node only. This information is not propagated to other nodes; you must copy this custom volume group processing method to each node manually. Alternatively, you can use the PowerHA SystemMirror File Collections facility to make the disk, volume, and file system methods available on all nodes.

Once you have made the selections, the information is applied to the volume groups on the local node.

4. Configure the same custom volume group type on other nodes in the cluster and synchronize the cluster resources. You may configure it manually or use the PowerHA SystemMirror File Collections facility. The cluster verification process ensures that the type that you configured exists and is executable on all nodes. PowerHA SystemMirror also verifies the root user permissions for the type

and methods you specified, and the fact that the methods do not reside on a shared physical volume. The synchronization process ensures that the entries in the PowerHA SystemMirror configuration database are the same on all nodes and synchronizes the methods named in the PowerHA SystemMirror file collections. PowerHA SystemMirror runs the specific verification checks for the OEM-type volume groups only if they are configured in the cluster.

Note: After you configure Veritas volume groups or file systems in PowerHA SystemMirror cluster, the PowerHA SystemMirror cluster verification utility runs the associated Veritas verification check. This verification check is effective only for the volume groups or file systems that are currently online on the node on which the verification is being run.

Changing custom volume group methods:

You can change custom volume group methods using SMIT.

To change a custom volume group method, complete the following steps:

1. Enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Resources > Custom Volume Group Methods > Add Custom Volume Group Methods** and press Enter. SMIT displays a picklist containing the names of the specified volume group processing types (or methods).
3. Select a name of a particular volume group type and press Enter. SMIT displays the current information.
4. Enter new information in the fields you want to change and press Enter.

Note: Unless you use the PowerHA SystemMirror file collections function, the method's characteristics are changed on the local node, but they are not updated on remote nodes. Manually update the custom volume group method on other nodes in the cluster.

Removing custom volume group methods:

You can remove custom volume group methods using SMIT.

To remove a custom volume method, complete the following steps:

1. Enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Resources > Custom Volume Group Methods > Remove Custom Volume Group Methods**, and press Enter.
SMIT displays a picklist with the names of the specified volume group types (or processing methods).
3. Select a type or method you want to remove and press Enter. SMIT prompts you to confirm your selection. If you choose to continue, PowerHA SystemMirror will delete the corresponding entry.

Note: Unless you use the PowerHA SystemMirror file collections function, PowerHA SystemMirror deletes the characteristics of the volume group types or methods on the local node but does not update them on remote nodes. Manually delete custom types or methods on other nodes in the cluster.

Integrating OEM file systems in a PowerHA SystemMirror cluster

You can configure OEM file systems in the AIX operation system and use PowerHA SystemMirror as an IBM High Availability solution to manage such volume groups, their corresponding file systems, and application controllers.

Application controllers are defined as start and stop scripts for the applications supported by the volume groups.

Overview for integrating OEM file systems

Depending on the type of OEM volume, custom methods in PowerHA SystemMirror allow you (or an OEM vendor) to tell PowerHA SystemMirror that a file system unknown to AIX LVM should be treated the same way as a known and supported file system, or to specify the custom methods that provide the file systems processing functions supported by PowerHA SystemMirror.

You can either use custom methods for file systems offered by PowerHA SystemMirror, or create your own custom methods to use for the non-IBM file systems in the cluster. These functions are performed within the normal PowerHA SystemMirror event processing.

In particular, PowerHA SystemMirror automatically detects and provides the methods for file systems created with Veritas File System (VxFS) using Veritas Foundation Suite v. 4.0.

Prerequisites and limitations:

In order for PowerHA SystemMirror to handle OEM volume groups and file systems in your configuration, the volume groups and file systems must include these prerequisites. There also are some limitations.

The volume groups and file systems must adhere to these conditions:

- Each OEM volume is composed of one or more physical disks (equivalent to `hdisk`s in AIX) and the system can determine the names of the physical disks from the name of the volume.
- OEM disks, volume groups, and file systems must have operations and sequences of operations comparable to the functions of AIX LVM, although the command names, syntax, and arguments may differ.

Limitations

PowerHA SystemMirror identifies OEM volume groups and file systems and performs all major functions with them.

However, some of the PowerHA SystemMirror functions have limitations for OEM volume groups and file systems:

- You cannot use C-SPOC (Cluster Single Point of Control) to manage OEM disk, volume, and file systems operations.
- PowerHA SystemMirror does not automatically discover OEM volume groups and file systems and does not list them for your selections in picklists in SMIT.
- PowerHA SystemMirror does not use NFS (Network File System) functions for OEM file systems.
- PowerHA SystemMirror does not provide a workaround for any limitations of OEM disks, volume groups, and file systems that exist outside of PowerHA SystemMirror.
- In addition to listing, varying on and off, and verifying volume groups and file systems created with the AIX LVM, PowerHA SystemMirror supports other extended functions for its "own" volume groups and file systems, such as enhanced concurrent mode, active and passive varyon process, disk heartbeating and selective fallover upon volume group loss and other.

These functions of PowerHA SystemMirror utilize the AIX LVM capabilities, and since OEM volume groups and file systems do not use AIX LVM, PowerHA SystemMirror does not support these functions for OEM volume groups.

- The Automatic Error Notification methods available in AIX cannot be configured in PowerHA SystemMirror for OEM volume groups and file systems.
- The PowerHA SystemMirror's mechanism for fast disk takeover cannot be utilized for mixed physical volumes, that is, for disks comprised of both IBM and non-IBM disks. The LVM enhanced concurrent mode capabilities of AIX are used to enable fast disk takeover - only LVM supported disks can be used with enhanced concurrent mode and the Fast Disk Takeover feature.

- PowerHA SystemMirror uses its serial method of processing resource groups if they contain OEM volume groups or file systems.
- You can define sites in the cluster, but you cannot configure and use replicated resources for PowerHA SystemMirror Enterprise Edition in the cluster that uses OEM volume groups. The intersite management policy for resource groups must be set to Ignore.
- PowerHA SystemMirror does not support the LVM cross-site mirroring function for resource groups that have OEM volume groups and file systems.

Software requirements

PowerHA SystemMirror supports Veritas volume groups at whatever level of VFS software is running on the system (depending on the version level supported by VFS 4.0).

OEM file systems functions performed by PowerHA SystemMirror:

When PowerHA SystemMirror identifies OEM file systems of a particular type, it provides the processing functions for them or you can specify your own custom methods for any one of these functions.

These processing functions include:

- Determine a list of file systems belonging to a specified type of volume group.
- List volume groups hosting a specified file system.
- Bring the file system online and offline.
- Determine the file system status.
- Verify the file systems configuration.
- Provide the path names to the file systems log files, for troubleshooting purposes.

OEM volume groups and file systems as resources in a PowerHA SystemMirror resource group:

You can include OEM disks, volume groups, and file systems in a PowerHA SystemMirror resource group. PowerHA SystemMirror recognizes OEM volume groups and file systems and handles them as the AIX LVM volume groups and file systems by listing, activating, checking and taking them offline when it is necessary for the cluster event processing tasks.

You can mix the OEM disks, volume groups and file systems with AIX volume groups and file systems in a single PowerHA SystemMirror resource group.

When you include OEM or Veritas volume groups or file systems into PowerHA SystemMirror resource groups, **Automatically Import Volume Groups** field should be set to **False** . Also OEM volume groups must be set to not automatically varyon when the node is restarted. OEM file systems must be set to not automatically mount when the node is restarted.

To view OEM disks, volume groups and file systems used in the PowerHA SystemMirror cluster, you can use the PowerHA SystemMirror SMIT interface.

Supporting Veritas file systems

Among other OEM volume groups and file systems, PowerHA SystemMirror supports volume groups and file systems created with VxVM in Veritas Foundation Suite v.4.0.

To make it easier for you to accommodate Veritas file systems in the PowerHA SystemMirror cluster, the methods for Veritas file systems support are predefined in PowerHA SystemMirror. After you add Veritas file systems to PowerHA SystemMirror resource groups, you can select the methods for the file systems from the picklists in PowerHA SystemMirror SMIT menus for OEM file systems support.

Note: After you configure Veritas volume groups or file systems in PowerHA SystemMirror cluster, the PowerHA SystemMirror cluster verification utility runs the associated Veritas verification check. This verification check is effective only for the volume groups or file systems that are currently online on the node on which the verification is being run.

Configuring OEM file systems methods in SMIT

You can add, change, and remove custom file systems processing methods for a specific OEM file system using SMIT. You can select existing custom file systems methods that are supported by PowerHA SystemMirror, or you can use your own custom methods.

Adding custom file systems methods:

You can add custom file systems methods using SMIT.

To add a custom file system method:

1. Enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Resources > Custom File System Methods > Add Custom File System Methods**, and press Enter.

Note: If you are configuring Veritas file systems, perform this operation in SMIT on the same node on which the file system is currently active.

3. Enter field values that define the file system processing methods you want to specify for the particular OEM file system type:

Table 26. Add Custom File System Method fields

Field	Value
File system Type	Enter the identifier for the particular file system type for which you want to configure existing methods in PowerHA SystemMirror or use your own custom methods. By default, it is the value of the VFS field of the <code>/etc/filesystems</code> file for the volume. Use the lsfs command to obtain this value from the <code>/etc/filesystems</code> file.
Method for listing file system names	Select a method from the picklist, or enter the full path name of a custom method that PowerHA SystemMirror must use to list file system names for the specified volume type. As a default method, PowerHA SystemMirror uses LSFS which calls the AIX lsfs command. The custom method that you use must generate a list of the file system names as an output. For PowerHA SystemMirror, the method returns zero for success and nonzero for failure. If a list is passed to PowerHA SystemMirror, the list must have one item per line.
Method for listing volume groups hosting a specified file system	Select a method from the picklist, or enter the full path name of a custom method that PowerHA SystemMirror must use to list volume groups that host the file system. As a default method, PowerHA SystemMirror compares the AIX logical volume that is associated with the file system with the AIX logical volume that is part of the volume group. By default, PowerHA SystemMirror uses ODM, which calls the <code>/usr/es/sbin/cluster/events/utlis/cl_fs2disk</code> command. The custom method that you use must generate a list of the file system names as an output. For PowerHA SystemMirror, the method returns zero for success and nonzero for failure. If a list is passed to PowerHA SystemMirror, the list must have one item per line.

Table 26. Add Custom File System Method fields (continued)

Field	Value
Method for bringing the file system online	<p>Select a method from the picklist, or enter the full path name of a custom method that PowerHA SystemMirror must use to activate the file system.</p> <p>As a default method, PowerHA SystemMirror uses MOUNT which calls the AIX mount command.</p> <p>For PowerHA SystemMirror, the method returns zero for success and nonzero for failure.</p>
Method for bringing the file system offline	<p>Select a method from the picklist, or enter the full path name of a custom method that PowerHA SystemMirror must use to take the file system offline.</p> <p>As a default method, PowerHA SystemMirror uses UMOUNT, which calls the AIX umount command.</p> <p>For PowerHA SystemMirror, the method returns zero for success and nonzero for failure.</p>
Method for determining the files ystem status	<p>Select the method from the picklist, or enter the full path name of a custom method that PowerHA SystemMirror must use to determine the file system status.</p> <p>As a default method, PowerHA SystemMirror uses LSACTIVEFS which calls the mount 2>/dev/null awk \$3 ~ /jfs2*\$/ {print \$2} command.</p> <p>For PowerHA SystemMirror, the method returns 0 if the file system is not mounted, 1 for the file system that is mounted, and 2 for command failure.</p>
Method for verifying file system configuration	<p>Enter the full path name of a custom method that PowerHA SystemMirror should use to verify the file system configuration on each cluster node.</p> <p>Note: By default, PowerHA SystemMirror verifies the configuration for AIX volume groups and file systems. It also uses the Veritas verification method for Veritas volume groups and file systems. However, the Veritas verification is effective only for those volume groups and file systems that are currently active on the nodes on which the verification is being run. For PowerHA SystemMirror, the method returns zero for success and nonzero for failure.</p>
Directories containing log information (optional)	<p>Enter one or more space-separated path names to the log files in this field or leave the field blank. If PowerHA SystemMirror finds the files in this field, you can use the snmp -e command to troubleshoot the log files.</p>

Note: The custom file system processing method or custom file system type is added to the local node only. This information is not propagated to other nodes. You copy this method or type to each node manually. Alternatively, you can use the PowerHA SystemMirror File Collections facility to make the disk, volume, and file system methods available on all nodes.

4. After you make the selections, the information is applied to the file systems on the local node.
5. Configure the same custom file systems processing method on other nodes in the cluster and synchronize the cluster resources. The cluster verification process ensures that the method (or the file system type) that you configured exists and is executable on all nodes. PowerHA SystemMirror also verifies the root user permissions for the methods you specified and the fact that the methods do not reside on a shared physical volume. The synchronization process ensures that the entries in the PowerHA SystemMirror configuration database are the same on all nodes and synchronizes the methods named in the PowerHA SystemMirror file collections.

Note: After you configure Veritas volume groups or file systems in PowerHA SystemMirror cluster, the PowerHA SystemMirror cluster verification utility runs the associated Veritas verification check. This verification check is effective only for the volume groups or file systems that are currently online on the node on which the verification is being run.

Changing custom file systems methods:

You can change custom file systems methods using SMIT.

To change a custom file system method:

1. Enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Resources > Custom File System Methods > Change/Show Custom File System Methods** and press Enter.
SMIT displays a picklist with the names of the specified file systems processing methods (or types).
3. Select a name of a particular file system method (or type) and press Enter. SMIT displays the current information.
4. Enter new information in the fields you want to change and press Enter.

Note: Unless you use the PowerHA SystemMirror file collections function, the type or method characteristics are changed on the local node but they are not updated on remote nodes. Manually update the custom methods or types on other nodes in the cluster.

Removing custom file systems methods:

You can remove custom file systems methods using SMIT.

To remove a custom file system method, complete the following steps:

1. Enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Resources > Custom File System Methods > Remove Custom File System Methods** and press Enter.
SMIT displays a picklist with the names of the specified file system types or processing methods.
3. Select the file system type or method you want to remove, and press Enter. SMIT prompts you to confirm your selection. If you choose to continue, PowerHA SystemMirror will delete the corresponding entry.

Note: Unless you use the PowerHA SystemMirror file collections function, PowerHA SystemMirror deletes the characteristics of the types or methods on the local node but does not update them on remote nodes. Manually delete custom types or methods on other nodes in the cluster.

PowerHA SystemMirror and SNMP utilities

These topics discuss the Simple Network Management Protocol (SNMP) and describes the relationship between the PowerHA SystemMirror SNMP-based utilities and other SNMP-based utilities that run on POWER® processor-based systems. This section does not discuss the SNMP standard in depth.

SNMP is a set of standards for monitoring and managing TCP/IP-based networks. SNMP includes a protocol, a database specification, and a set of data objects. A set of data objects forms a Management Information Base (MIB). SNMP provides a standard MIB that includes information such as IP addresses and the number of active TCP connections. The actual MIB definitions are encoded into the agents running on a system. The standard SNMP agent is the SNMP daemon, **snmpd**.

MIB-2 is the standard for defining over 100 TCP/IP specific objects, including configuration and statistical information such as:

- Information about interfaces
- Address translation
- IP, ICMP (Internet-control message protocol), TCP, and UDP

SNMP can be extended through the use of the SNMP Multiplexing protocol (the SMUX protocol) to include enterprise-specific MIBs that contain information relating to a discrete environment or application. The Cluster Manager retrieves and maintains information about the objects defined in its MIB, and passes this information on to a specialized network monitor or network management station.

The PowerHA SystemMirror software, NetView® for AIX, and Systems Monitor for AIX all include the following SMUX daemons: **clstrmgr**, **trapgend**, and **sysinfod**, respectively. You must be aware of possible conflicts between these daemons.

PowerHA SystemMirror SNMP components

The PowerHA SystemMirror software provides an enterprise-specific (generic type 6) MIB. The source file is **hacmp.my**. The **mosy** command compiles **hacmp.my** (with other standard MIBs) to generate the **hacmp.defs** file.

The PowerHA SystemMirror MIB, is associated with and maintained by the Cluster Manager. The PowerHA SystemMirror software also provides two cluster monitor programs, the Cluster Information Program (Cinfo) and **clstat**.

Cluster information (Cinfo) program

Cinfo is a cluster monitoring program. It requests information about the current cluster state from the Cluster Manager. The Cluster Manager updates data using internal, dynamically allocated data structures that are accessible to Cinfo clients applications that use Cinfo API functions.

By default, Cinfo receives information from the Cluster Manager by polling. The time between polling is set by an argument to Cinfo, which defaults to 15. Cinfo can also receive information asynchronously through traps. In response to traps, Cinfo sends a request for more information to the Cluster Manager. It does not parse the trap message data itself; instead, Cinfo employs a trap-directed polling policy.

To enable Cinfo to receive traps, call it using the **-a** option. Since Cinfo is started through the System Resource Controller (SRC), the best way to do this is by entering:

```
chssys -s cinfoES -a "-a"
```

Then use the **lssrc** command to ensure this change occurred. Enter:

```
lssrc -Ss cinfoES | awk -F: '{print $3}'
```

Traps provide more timely information to Cinfo clients. The clients simply register to receive various events. Cinfo notifies the traps by using signals when those events occur. However, that Cinfo's polling interval is doubled when traps are enabled.

SNMP community names and Cinfo

The default SNMP community name for Cinfo is public. You can override this by using the following command to force the SRC to start Cinfo with the **-c** switch by entering:

```
chssys -s cinfoES -a "-c abcdef"
```

here abcdef is an SNMP community name defined as such in the **snmpd** configuration file.

Then use the **lssrc** command to ensure this change occurred. Enter:

```
lssrc -Ss cinfoES | awk -F: '{print $3}'
```

The Simple Network Management Protocol (SNMP) community name used by PowerHA SystemMirror depends on the version of SNMP you are running on your system. The SNMP community name is determined as follows:

- If your system is running SNMP V1, the community name is the first name found that is not private or system in the output of the **lssrc -ls snmpd** command.

- If your system is running SNMP V3, the community name is found in the VACM_GROUP entry in the `/etc/snmpdv3.conf` file.

The Clinfo service still supports the `-c` option for specifying SNMP community name but its use is *not* required. The use of the `-c` option is considered a security risk because doing a `ps` command could find the SNMP community name. If it is important to keep the SNMP community name protected, change permissions on `/tmp/hacmp.out`, `/etc/snmpd.conf`, `/smit.log` and `/usr/tmp/snmpd.log` to *not* be world readable.

Important notes on snmpdv3.conf file

AIX `snmpdv3` has three functions or parts: One is the SNMP v3 agent, one is the SMUX server, and the last is the DPI2 agent. The DPI2 agent has to use community "public" to get a port number from DPI2 subagents (`hostmibd` , `snmpmibd` , `aixmibd`) to communicate with them. For this reason you should still keep the community name as public and give the public a view of only this `dpiPortForTCP.0` (1.3.6.1.4.1.2.2.1.1.0) MIB variable so that the DPI2 agent can get the port number from subagents. See the example `snmpdv3.conf` file below.

Sample snmpdv3.conf file with nonpublic community name

```
VACM_GROUP group1 SNMPv1 YourLongNameHere -
VACM_GROUP group2 SNMPv1 public -

VACM_VIEW defaultView internet - included -
VACM_VIEW dpi2view 1.3.6.1.4.1.2.2.1.1.1 - included -

VACM_VIEW defaultView snmpModules - excluded -
VACM_VIEW defaultView 1.3.6.1.6.3.1.1.4- included -
VACM_VIEW defaultView 1.3.6.1.6.3.1.1.5- included -

VACM_VIEW defaultView 1.3.6.1.4.1.2.6.191 - excluded -

VACM_ACCESS group1 - - noAuthNoPriv SNMPv1 defaultView - defaultView -
VACM_ACCESS group2 - - noAuthNoPriv SNMPv1 dpi2view - - -

NOTIFY notify1 traptag trap -

TARGET_ADDRESS Target1 UDP 127.0.0.1traptag trapparms1 - - -

TARGET_PARAMETERS trapparms1 SNMPv1 SNMPv1 YourLongNameHere
noAuthNoPriv -

COMMUNITY YourLongNameHere YourLongNameHere noAuthNoPriv 0.0.0.0
0.0.0.0 -
COMMUNITY public public noAuthNoPriv 0.0.0.0 0.0.0.0 -

DEFAULT_SECURITY no-access - -

logging file=/usr/tmp/snmpdv3.logenabld
logging size=4194304level=0

smux 1.3.6.1.4.1.2.3.1.2.1.2 gated_password # gated

smux 1.3.6.1.4.1.2.3.1.2.3.1.1 muxatmd_password #muxatmd
smux 1.3.6.1.4.1.2.3.1.2.1.5clsmuxpd_password #PowerHA SystemMirror/ES for
AIX clsmuxpd
```

Related information:

`snmpd.conf` file

SNMP for network management

The `/usr/es/sbin/cluster/clstat` utility

The `/usr/es/sbin/cluster/clstat` utility runs on both ASCII and X terminals. The display automatically corresponds to the capability of the system. However, if you want to run an ASCII display on an X-capable machine, you can do so by specifying the `-a` option.

In addition, you can set up `clstat` to display in a web browser, if you set up a web server on a node that has `clinfo` running. The browser display makes it easier to view multiple clusters without having to select the clusters one at a time.

`clstat` is a Clinfo client. It uses the Clinfo C API to get cluster information from the shared memory segment maintained by Clinfo. It does not register to receive events, but uses the Clinfo polling method.

The LPP contains both executables and source code for the `clstat` utility. If you want to recompile `clstat`, run the `make` command in the directory `/usr/es/sbin/cluster/samples/clstat`.

Systems monitor for AIX

Systems monitor for the AIX operating system runs the `sysinfod` SNMP multiplexing (SMUX) peer daemon that monitors certain characteristics.

These characteristics include:

- Machine name, type, and processor ID
- Devices installed on the machine
- Operating system configuration
- Status of subsystems, paging devices, and file systems
- Network traffic
- Active processes
- Users
- CPU and device utilization

If trap filtering is enabled on this agent system, the `sysinfod` daemon receives SNMP traps on port 162. By default, the `snmpd` daemon sends all received SNMP traps to the `sysinfod` daemon for filtering. The traps are evaluated by the `sysinfod` daemon, and those traps meeting the filter criteria are forwarded to the manager system.

Systems monitor startup options for PowerHA SystemMirror compatibility

If you are using the Systems Monitor for AIX along with PowerHA SystemMirror on your system, start the `sysinfod` with the `-H` option. This option allows the PowerHA SystemMirror `cl_swap_HW_address` utility to function correctly. If the `sysinfod` is not started with the `-H` option, it keeps the adapter busy all the time it is active, and this prevents the `cl_swap_HW_address` utility from removing the device when it tries swapping the HW address.

Trap conflicts between SMUX peer daemons

A single SNMP agent (`snmpd` daemon) can send the same trap to multiple SNMP managers; this agent is configured in the `snmpd` configuration file. However, only one SNMP manager (for example, NetView for AIX) can run on a given network station, because only one TCP/IP program at a time can listen on a particular port. There is no way to work around this limitation.

In the case of NetView for AIX, the `trapd` daemon listens on port 162 and forwards traps to NetView for AIX. In turn, NetView for AIX can forward traps to multiple NetView for AIX applications that have registered with NetView for AIX. The `trapgend` daemon can generate traps for AIX system error-log-related events. The variables in the private portion of `trapgend` are described in the file `/usr/etc/nm/mibs/ibm-nv6ksubagent.mib`.

When the **sysinfod** daemon is installed on an NetView for AIX manager, trap reception is disabled for filtering. This is set in the **/usr/adm/sm6000/config/install.config** configuration file. However, when the **sysinfod** daemon is installed on a node without the manager installed, trap reception is enabled using the same file. You can install NetView for AIX on a node where the **sysinfod** daemon is already installed and trap reception is enabled. This causes the NetView for AIX **trapd** daemon to fail to start since the **sysinfod** daemon is using the port.

Both the NetView for AIX manager and the **sysinfod** daemon cannot share this port. Disable filtering on this node by way of the **/usr/adm/sm6000/config/install.config** configuration file. In this way, when you start the **sysinfod** daemon, it has trap reception and filtering disabled.

Similarly, Clinfo cannot be enabled to receive traps from the SNMP process (activated by the **-a** flag) if **trapgend** is also running. If the NetView for AIX trap daemons are started first, Clinfo will immediately exit with a `smux_connect` error. If Clinfo is started first with the **-a** option, most of the NetView for AIX daemons will *not* start.

Notices

This information was developed for products and services offered in the US.

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

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

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

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

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

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

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

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

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

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

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

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

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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

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

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 actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

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

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

© (your company name) (year).

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

© Copyright IBM Corp. _enter the year or years_.

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.

Index

A

- adding
 - copies to logical volume 37
- AIX error notification function 49

B

- breaking
 - disk reserve 58

C

- CD-ROM,
 - installing on server node 28
- changing
 - volume group
 - startup status 40
- client node
 - installing PowerHA SystemMirror 31
- clinfo 72
- clstat 74
- cluster
 - configuring
 - OEM volume group 64
 - configuring in SMIT
 - OEM file systems 69
 - integrating
 - OEM disks 55
 - OEM file systems 66
 - OEM volume group 61
 - SNMP
 - clinfo 72
 - clstat 74
 - components 72
 - overview 71
- configuring
 - network interface cards 33
 - OEM disks
 - SMIT 59
 - OEM file systems
 - SMIT 69
 - OEM volume group 64
 - tape drive
 - shared 34
 - to use NFS 4
- creating
 - shared file system 36
 - shared LVM components 2
 - shared volume group 36
- customizing
 - disk processing 57

D

- defining
 - LVM components
 - concurrent access 41
 - shared LVM components 35
 - for nonconcurrent 35

- disk
 - breaking reserve 58
 - customizing processing 57

E

- error notification
 - AIX 49
 - PowerHA SystemMirror 50
 - SIGDANGER 48
 - test 54

F

- file collection
 - planning 46
- file system
 - creating shared 36
 - testing 37

H

- hard disk
 - installing on server node 27

I

- importing
 - volume group 38
- installation server
 - installing on server node 27
- installing
 - on client nodes 31
 - on server nodes 23
 - CD-ROM 28
 - completing install 30
 - from hard disk 27
 - installation server 27
 - prerequisites 23
 - problems 30
 - tape drive
 - shared Fibre 33
- integrating
 - OEM disks
 - cluster 55
 - OEM file systems 66
 - overview 67
 - OEM volume group 61
 - overview 61

J

- jfslog
 - renaming 36

L

- logical partition 3
- logical volume 3
 - adding copies 37
 - renaming 36
- LVM component
 - defining for concurrent access 41

M

- migration
 - troubleshooting 22
 - verifying 20

N

- network considerations 45
- network interface card
 - configuring 33
- NFS
 - configuring to use 4
- Non-Disruptive Upgrade 15

O

- OEM disk
 - configuring
 - SMIT 59
 - integrating in cluster 55
- OEM file systems
 - configuring in SMIT 69
 - integrating in cluster 66
 - overview 67
 - Veritas 68
- OEM volume group
 - configuring in cluster 64
 - integrating in cluster 61
 - overview 61
 - Veritas 64
- original equipment manufacturer
 - See OEM
- overview 1
 - integrating
 - OEM file systems 67
 - OEM volume group 61
 - SNMP 71
 - system monitor 74

P

- planning
 - file collections 46
 - for upgrade 10
- PowerHA SystemMirror automatic error notification facility 50

R

- recovering
 - previous version of PowerHA SystemMirror 22
- removing
 - from server nodes
 - PowerHA SystemMirror software 30

- renaming
 - jfslogs 36
 - logical volumes 36

S

- server node
 - completing install 30
 - installing
 - from CD-ROM 28
 - from hard disk 27
 - from installation server 27
 - installing PowerHA SystemMirror 23
 - problems during installation 30
 - removing PowerHA SystemMirror software 30
- shared LVM component
 - creating 2
 - defining 35
 - for nonconcurrent 35
- Simple Network Management Protocol
 - See SNMP
- snapshot
 - upgrade from 7.1.3, or later 16
 - upgrade using 16
- SNMP
 - clinfo 72
 - clstat 74
 - components 72
 - overview 71
- system monitor
 - overview 74

T

- tape drive
 - configuring
 - shared 34
 - installing
 - shared Fibre 33
- testing
 - file system 37
- troubleshooting
 - migration 22

U

- upgrading
 - migration 12
 - networks 11
 - offline cluster 17
 - offline cluster from 7.1.3, or later 17
 - planning for upgrade 10
 - prerequisites 9
 - rolling migration 14
 - rolling migration from 7.1.3 to 7.2.0 or later 14
 - security 12
 - using a snapshot 16
 - using a snapshot from PowerHA SystemMirror 7.1.1, or later 16

V

- varying off
 - volume group 38, 40

- verifying
 - migration 20
- Veritas file systems 68
- Veritas volume groups 64
- volume group
 - changing startup status 40
 - creating shared 36
 - importing 38
 - varying off 38, 40



Printed in USA