

**IBM Geographic Logical
Volume Manager Geographic
Logical Volume
Manager Geographic Logical
Volume Manager**

Contents

| | |
|----------------------------------------------------------------------------------------------------------|----------|
| Highlighting..... | 6 |
| Case-sensitivity in AIX® | 6 |
| ISO 9000..... | 6 |
| Related information | 6 |
| Geographic Logical Volume Manager for PowerHA® SystemMirror® Enterprise Edition.... | 7 |
| What's new in Geographic Logical Volume Manager (GLVM) for PowerHA® SystemMirror® Enterprise Edition .. | 7 |
| How to see what's new or changed..... | 7 |
| December 2023..... | 7 |
| December 2022..... | 7 |
| December 2021..... | 7 |
| Geographic Logical Volume Manager for PowerHA® SystemMirror® Enterprise Edition concepts..... | 7 |
| Data mirroring..... | 8 |
| Geographic Logical Volume Manager (GLVM) | 10 |
| Remote physical volume (RPV) | 11 |
| Stand-alone GLVM technology | 12 |
| GLVM technology combined with the PowerHA® SystemMirror® automated monitoring and data | |
| recovery..... | 12 |
| High availability and disaster recovery support features | 13 |
| Synchronous mirroring..... | 13 |
| Asynchronous mirroring..... | 14 |
| GLVM for a PowerHA® SystemMirror® cluster | 15 |
| Example: From ordinary LVM mirroring to GLVM | 19 |
| Example: Node and site failure..... | 21 |
| Planning for GLVM..... | 23 |
| Planning geographically mirrored volume groups | 23 |
| Planning PowerHA® SystemMirror® clusters with geographically mirrored volume groups | 26 |
| Planning for asynchronous mirroring..... | 32 |
| Installing GLVM for PowerHA® SystemMirror® Enterprise Edition | 35 |
| Installation components | 35 |
| Installation prerequisites..... | 36 |
| Installation options..... | 36 |
| Filesets on the installation media..... | 37 |
| Installation methods | 37 |
| Addressing problems during the installation..... | 40 |
| Upgrading GLVM for PowerHA® SystemMirror® Enterprise Edition | 40 |
| Configuring geographically mirrored volume groups | 41 |
| Configuration prerequisites | 41 |
| GLVM Configuration Assistant | 42 |
| Prerequisites for using the GLVM Configuration Assistant | 46 |
| Configuring standard volume groups | 47 |
| Configuring enhanced concurrent mode volume groups | 47 |
| Configuration overview for asynchronous mirroring..... | 47 |
| Configuring RPV client and server pairs for a standard volume group at a node on each site | 48 |
| Extending the geographically mirrored standard volume group to other nodes in the cluster..... | 49 |
| Configuring RPV client-server pairs for an enhanced concurrent volume group on a node at each site | 49 |
| Extending an existing enhanced concurrent volume group to nodes at remote site | 49 |
| Configuring asynchronous mirroring..... | 50 |
| Configuring remote physical volume servers and clients..... | 52 |
| Configuring RPVs, mirror copies and geographically mirrored logical volumes | 56 |
| Using GLVM utilities: Managing RPVs and remote site mirror copies..... | 56 |
| Managing geographically mirrored volume groups with mirror pools..... | 60 |
| Adding a remote site mirror copy to a logical volume | 66 |
| Extending a logical volume that contains remote site mirrors..... | 67 |
| Removing a remote site mirror copy from a volume group..... | 68 |
| Listing all geographically mirrored logical volumes | 69 |
| Verifying mirror copy site locations for a logical volume | 69 |
| Changing the RPV servers and clients configuration | 70 |
| Changing RPVs and logical volume | 75 |
| Converting a standard cluster to a linked cluster..... | 78 |
| Integrating geographically mirrored volume groups into a PowerHA® SystemMirror® cluster | 80 |
| Identifying the configuration path..... | 80 |
| Geographically mirrored volume groups configuration prerequisites..... | 80 |
| Overview of configuration tasks..... | 81 |
| Configuring sites | 81 |
| Configuring XD-type networks..... | 81 |
| Configuring resource groups in GLVM for PowerHA® SystemMirror® Enterprise Edition | 83 |
| Verifying and synchronizing the GLVM configuration in PowerHA® SystemMirror® | 89 |
| Changing GLVM-mirrored resources configuration..... | 90 |

| | |
|-----------------------------------------------------------------------------------------------|------------|
| Managing GLVM for PowerHA® SystemMirror® Enterprise Edition..... | 91 |
| Managing RPVs outside of PowerHA® SystemMirror® control..... | 91 |
| Resuming the RPV clients manually..... | 91 |
| Managing geographically mirrored volume groups outside of PowerHA® SystemMirror® control..... | 92 |
| Troubleshooting GLVM for PowerHA® SystemMirror® Enterprise Edition | 93 |
| Logging messages | 93 |
| Troubleshooting the GLVM Configuration Assistant | 93 |
| Cache device failure | 93 |
| Asynchronous I/O cache logical volume failure | 93 |
| Running corrective actions during verification | 95 |
| Quorum and forced varyon in GLVM for PowerHA® SystemMirror® Enterprise Edition..... | 95 |
| Recovering from a disaster while using asynchronous mirroring | 101 |
| Monitoring the state of your GLVM environment..... | 104 |
| Notices..... | 117 |
| Privacy policy considerations | 118 |
| Trademarks..... | 118 |
| Index | 119 |

Note

Before using this information and the product it supports, read the information in [“Notices” on page 117](#).

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

About this document

This document introduces the PowerHA® SystemMirror® Enterprise Edition for AIX® software. This information is also available on the documentation CD that is shipped with the operating system.

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 for AIX® PDF documents are available in the [PowerHA SystemMirror 7.2 PDFs](#) topic.
- The PowerHA® SystemMirror® Version 7.2 for AIX® release notes are available in the [PowerHA SystemMirror 7.2 release notes](#) topic.

Geographic Logical Volume Manager for PowerHA® SystemMirror® Enterprise Edition

GLVM for PowerHA® SystemMirror® Enterprise Edition is only supported in PowerHA® SystemMirror® 7.1.2 Enterprise Edition for AIX®.

This information provides the necessary steps to plan, install, configure, and maintain Geographic Logical Volume Manager (GLVM) for PowerHA® SystemMirror® Enterprise Edition.

Note: PowerHA® SystemMirror® is the new name for HACMP™.

What's new in Geographic Logical Volume Manager (GLVM) for PowerHA® SystemMirror® Enterprise Edition

Read about new or significantly changed information for the Geographic Logical Volume Manager (GLVM) for PowerHA® SystemMirror® Enterprise Edition topic collection.

How to see what's new or changed

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

December 2023

Updated the following topic about GLVM disaster recovery (DR) Sizing tool that monitors data replication and provides suitable network latency recommendations using imperf tool:

- [GLVM DR Sizing tool](#)

December 2022

Added the following topic about GLVM disaster recovery (DR) Sizing tool that monitors data replication and provides suitable recommendations:

- [GLVM DR Sizing tool](#)

December 2021

The following topic is added to the [Configuring geographically mirrored volume groups](#) topic:

- [Converting a standard cluster to a linked cluster](#)

Updated a note about JSON file in the [GLVM Configuration Assistant enhancements](#) topic.

Geographic Logical Volume Manager for PowerHA® SystemMirror® Enterprise Edition concepts

Geographic Logical Volume Manager (GLVM) for PowerHA® SystemMirror® Enterprise Edition provides disaster recovery and data mirroring capability for the data at geographically separated sites. It protects the data against total site failure by remote mirroring, and supports unlimited distance between participating sites.

GLVM for PowerHA® SystemMirror® Enterprise Edition provides disaster recovery capability by exploiting the AIX® Logical Volume Manager and GLVM subsystems in order to create volume groups and logical volumes which span across two geographically separated sites.

You can use the GLVM technology as a standalone method, or use it in combination with PowerHA® SystemMirror® Enterprise Edition.

The software increases data availability by providing continuing service during hardware or software outages (or both), planned or unplanned, for a two-site cluster. The distance between sites may be unlimited, and both sites may access the mirrored volume groups serially over IP-based networks.

Also, it enables your business application to continue running at the takeover system at a remote site while the failed system is being recovered from a disaster or a planned outage.

The software takes advantage of the following software components to reduce downtime and recovery time during disaster recovery:

- AIX® LVM subsystem and GLVM
- TCP/IP subsystem
- PowerHA® SystemMirror® for AIX® cluster management

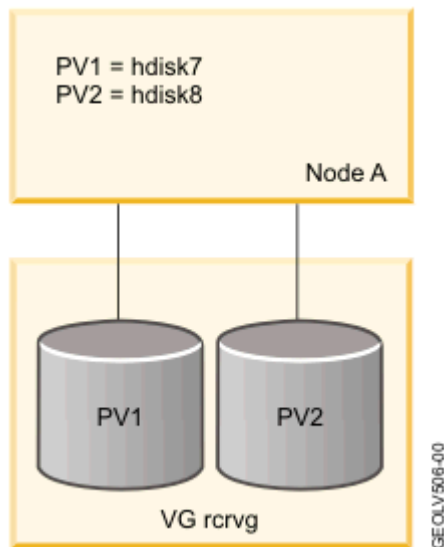
Data mirroring

Review this section to better understand the importance of data mirroring.

Configurations that allow users to access data can have different levels of protection against data loss. Consider the following data access configurations at a high level. All these configurations enable users to have access to data, with increasing levels of data protection against loss:

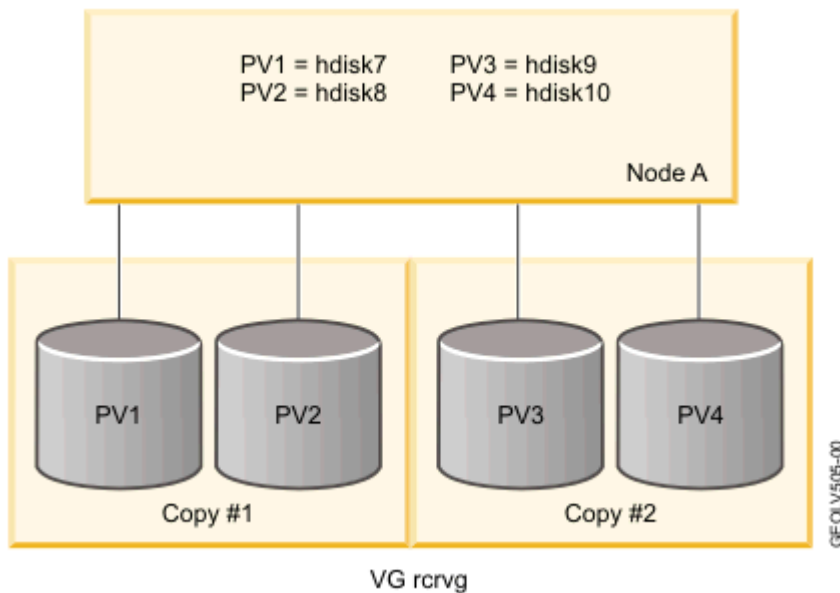
Single server with one data copy

In the following configuration, a single copy of mission-critical data spans two disks. Either disk is potentially a single point of failure:



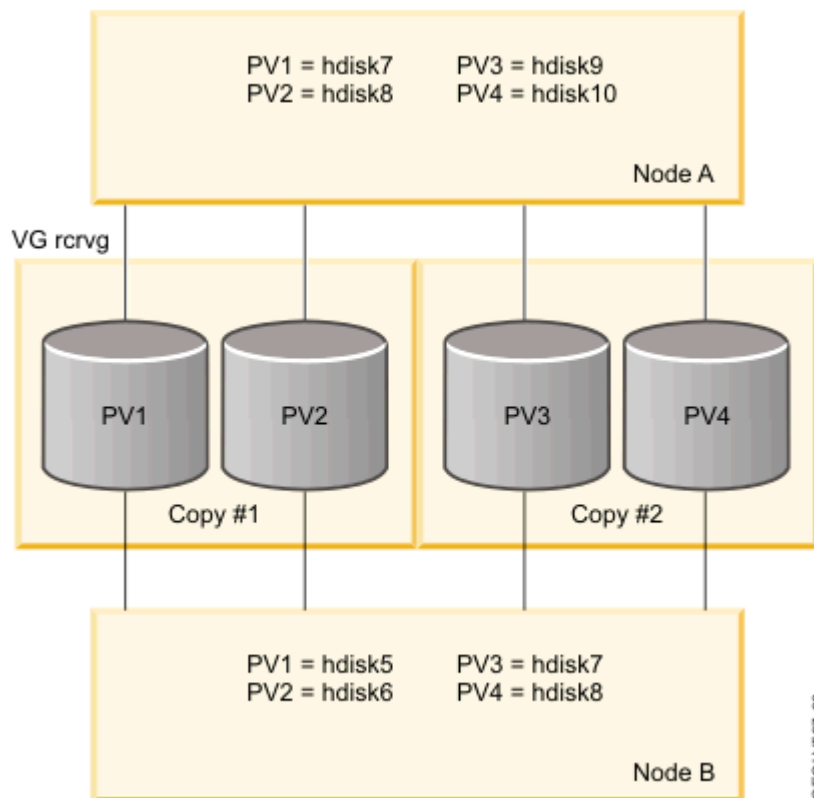
Single server with standard AIX® LVM mirroring

The following configuration has two more disks that contain a second mirror copy. However, the server is still a potential single point of failure.



PowerHA® SystemMirror® cluster with standard AIX® LVM mirroring

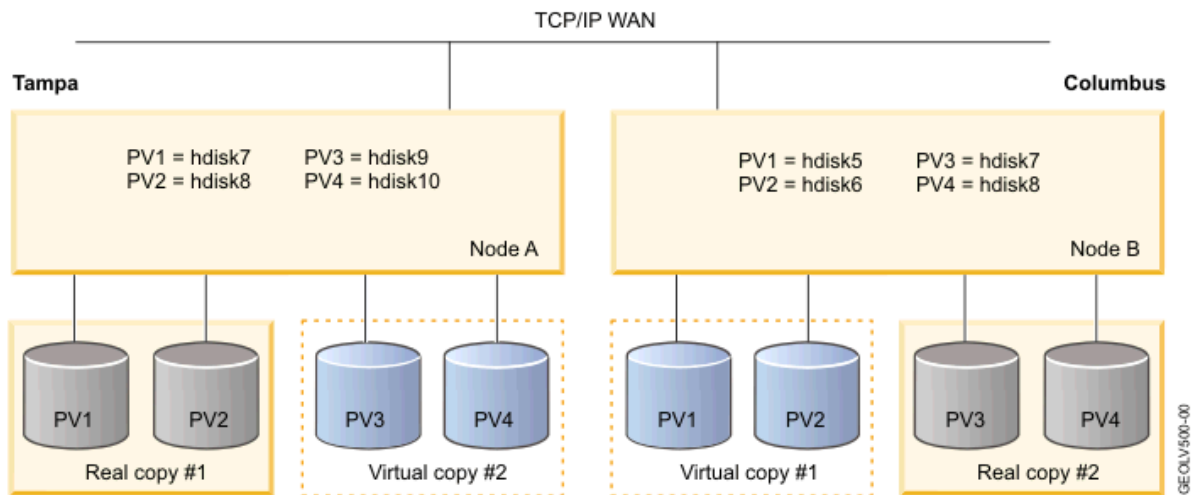
The following configuration is based on a PowerHA® SystemMirror® cluster that uses AIX® LVM mirroring:



This configuration eliminates most potential single points of failure. However, a disaster can wipe out the entire cluster along with all copies of mission critical data.

Two-site PowerHA® SystemMirror® cluster with GLVM mirroring

Finally, the configuration shown in the following image contains two sites with nodes at each site. One volume group actually spans both sites. Each site contains a mirror copy of mission critical data. Instead of extremely long disk cables, a TCP/IP network and the RPV device driver are used to enable remote disk access.



Geographic Logical Volume Manager (GLVM)

The Geographic Logical Volume Manager (GLVM) provides a software-based mirroring method in GLVM for PowerHA® SystemMirror® Enterprise Edition that is used for mirroring your data in real-time over unlimited geographic distance.

The main function of the GLVM is mirroring of any data destined for nodes at a local site (the production site) across an IP-based network to nodes at a remote site (the backup site). A total failure (such as a CPU, disk, network, or power failure) of the node at the local site does not cause the loss of data on the node at the remote site.

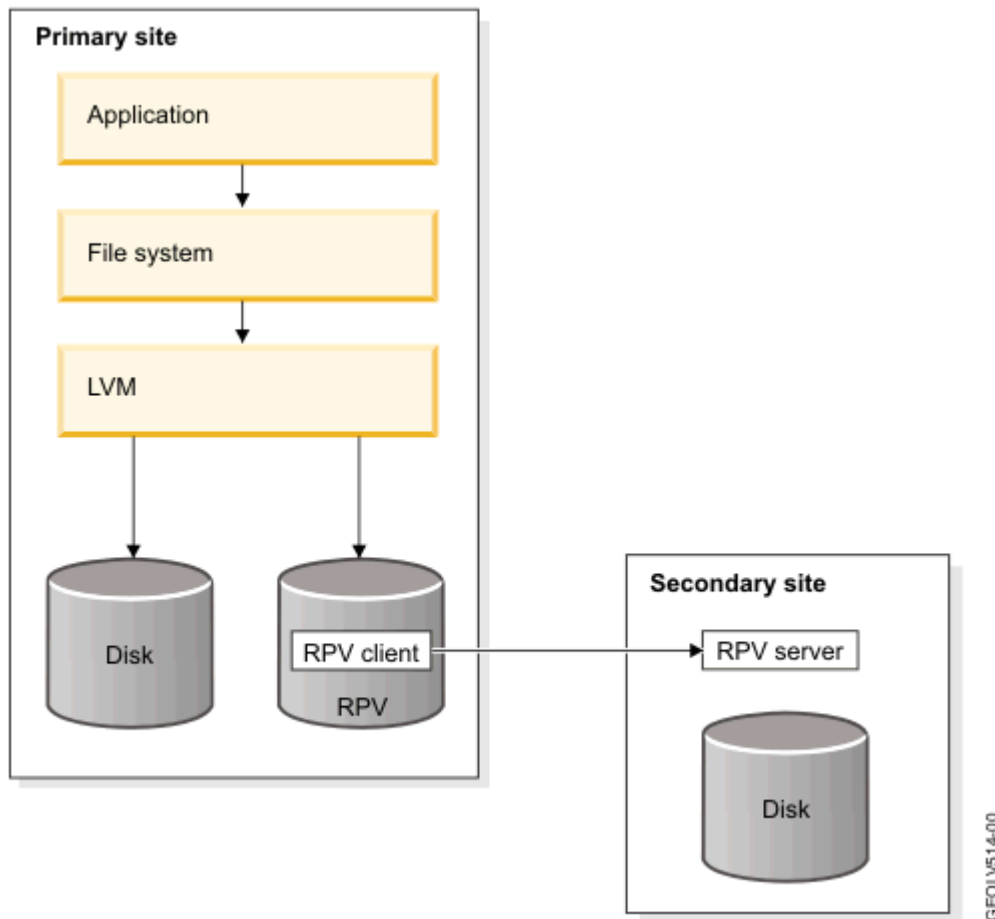
The GLVM mirroring:

- Does not depend on the type of database and file system. There is no modification required of applications that use GLVM mirroring capabilities.
- Performs the data mirroring over standard TCP/IP networks without having a dependency on the specific data that is being mirrored.
- Is less expensive than hardware-based mirroring solutions and does not have distance limitations.
- Uses the mirroring function of the AIX® LVM and operates as a layer under the LVM.

With the AIX® LVM and GLVM for PowerHA® SystemMirror® Enterprise Edition, you can create, change, and verify geographic mirror copies of your logical volumes. The mirror copies that you create must:

- Reside on different sites.
- Have the Super Strict inter-disk allocation policy for the logical volumes. This forces each complete copy of the logical volume to reside on separate sets of physical volumes and *remote physical volumes (RPVs)*.
- Not be part of the root volume group.

The following figure illustrates the GLVM for PowerHA® SystemMirror® Enterprise Edition mirroring function:



The Geographic LVM includes:

- The remote physical volume (RPV) device driver.
- You can use the SMIT interface to configure RPVs for a volume group that you plan to replicate remotely. You can also use the SMIT interface to change and verify mirror copies for logical volumes and volume groups.

The GLVM also uses the AIX® LVM mirroring capabilities.

Remote physical volume (RPV)

Remote physical volume (RPV) is a pseudo device driver and its kernel extension that allows the LVM to consider the physical volume located at the remote site as another local physical volume, although the actual data I/O operations are performed at the remote site.

You can mirror your data across two sites by configuring volume groups that contain both local and RPVs. With an RPV device driver, the LVM does not distinguish between local and remote physical volumes - it maintains mirror copies of the data across attached disks. The LVM is, for the most part, unaware that some disks are located at a remote site. During a disk read operation, however, the LVM tries to satisfy the read request from a local copy of the data (when available) to maximize performance. During a disk write operation the RPV device driver and LVM work together to allow the application to perform a single logical write that results in multiple physical writes to the local PVs and RPVs comprising the volume group.

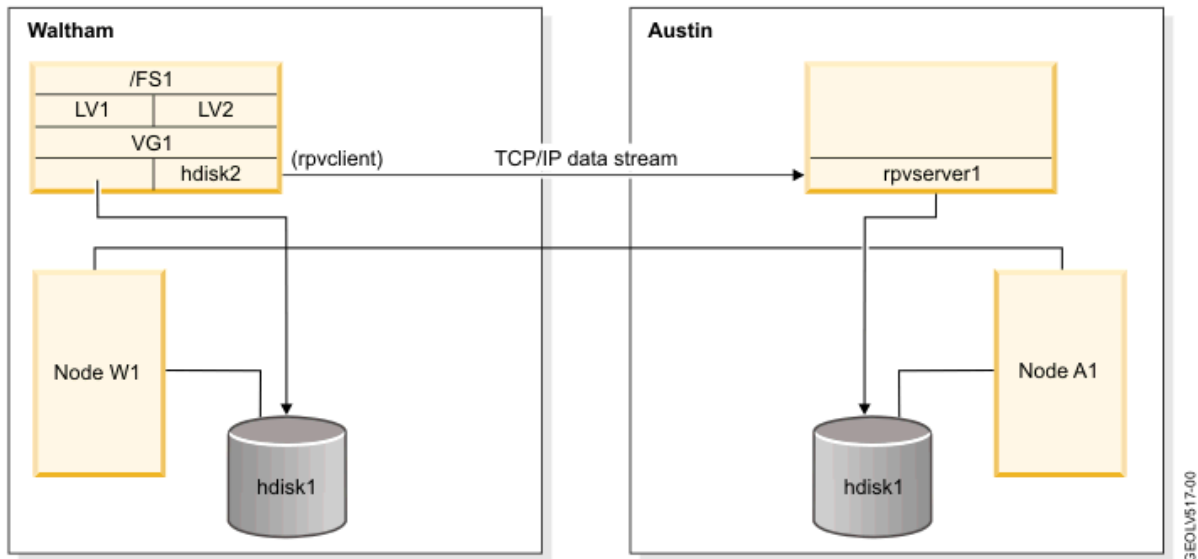
The RPV consists of:

- The *RPV client*. The RPV client is a pseudo device driver that runs on the node on the local site, that is, on the node at a geographic location where your applications issue data I/O requests to the physical volumes.
- The *RPV server*. The RPV server is an instance of the kernel extension of the RPV device driver that runs on the node on the remote site, that is, on the node at a remote geographic location to which the data is mirrored over an IP-based network. The RPV client sends I/O requests to the RPV server.

The geographically mirrored volume group is a volume group that contains one or more logical volumes that have copies residing on RPVs.

After you configure geographically mirrored volume groups which contain RPVs, you can include them as replicated resources in PowerHA® SystemMirror®-managed resource groups in the PowerHA® SystemMirror® cluster that is configured across two sites.

The following figure illustrates the mirroring process that occurs between the RPV client and RPV server. The RPV client is configured on the node at site Waltham (the local site for the resource group) and the RPV server is configured on the node at site Austin (the remote site for the resource group):



Stand-alone GLVM technology

You can configure geographically mirrored volume groups in AIX® GLVM, without having to install and configure a PowerHA® SystemMirror® cluster. The AIX® GLVM technology provides the same geographic data mirroring functionality as GLVM for PowerHA® SystemMirror® Enterprise Edition, only without the automated monitoring and recovery which is provided by PowerHA® SystemMirror®.

Consider using stand-alone AIX® GLVM technology to gain experience with remote mirroring before adding the fully automated failure detection and resource recovery provided by the PowerHA® SystemMirror® Enterprise Edition solution.

GLVM technology combined with the PowerHA® SystemMirror® automated monitoring and data recovery

GLVM provides a remote data mirroring function and a integration with PowerHA® SystemMirror® function.

Remote data mirroring

Geographic Logical Volume Manager (GLVM) for PowerHA® SystemMirror® Enterprise Edition creates a remote mirror copy of the data that the application can access both at the local site and at the site located at another, remote location.

The software protects critical data by mirroring the nonconcurrent and enhanced concurrent mode volume groups to which the application sends I/O requests. The application can access the same data regardless of whether the application is running on the local or remote site. (Note, the application can be running either on the local or on the remote site, but not at both sites at the same time.)

The GLVM mirroring technology is built upon the AIX® Logical Volume Manager subsystem. The same high availability concepts that are used to prevent disk failures from disrupting mission critical applications are expanded upon in order to keep a disaster which affects an entire computing center from disrupting mission critical applications on a larger scale.

The data mirroring function utilizes the mirroring capabilities of the AIX® Logical Volume Manager (LVM) along with the mirroring function of the GLVM that builds upon the LVM and is highly integrated with it.

Integration with PowerHA® SystemMirror®

By integrating with PowerHA® SystemMirror®, GLVM for PowerHA® SystemMirror® Enterprise Edition keeps mission-critical systems and applications operational in the event of disasters. GLVM for PowerHA® SystemMirror® Enterprise Edition manages the fallover and fallback of the resource group that contains the

application. The resource group is active at nodes on either site (with a nonconcurrent site management policy).

When a node, a network interface or a site fails, GLVM for PowerHA® SystemMirror® Enterprise Edition moves the resource group to another node. The node may belong either to the same site or to a remote site. With this operation, a complete up-to-date copy of the volume group's data remains available for the application to use on the node at the same site or at another site.

In GLVM for PowerHA® SystemMirror® Enterprise Edition, each site contains a mirror copy of a geographically mirrored volume group. If a disaster occurs at one site, the surviving site already has an up-to-date copy of the data. By including this data mirroring technology into a PowerHA® SystemMirror® cluster that also spans across both sites, PowerHA® SystemMirror® automates disaster recovery by detecting a site outage and then starting mission critical applications at the surviving site, which already has an up-to-date copy of your mission critical data.

GLVM for PowerHA® SystemMirror® Enterprise Edition protects against power outages, hardware or software failures, and natural disasters. This is accomplished by eliminating the system and the site as points of failure.

Note: PowerHA® SystemMirror® Standard Edition already protects against node failures. In addition to that, GLVM for PowerHA® SystemMirror® Enterprise Edition protects against site failures that can be caused by natural disasters, fires, terrorism and other unplanned calamities.

High availability and disaster recovery support features

GLVM for PowerHA® SystemMirror® Enterprise Edition has the several High Availability (HA) and disaster recovery support features.

These features include:

- Provides automatic detection and response to site and network failures in the geographic cluster without user intervention.
- Performs automatic site takeover and recovery and keeps mission-critical applications highly available through application failover and monitoring.
- Allows for simplified configuration of volume groups, logical volumes and resource groups. Supports either standard or enhanced concurrent volume groups that are geographically mirrored.
- Uses up to four TCP/IP networks for remote mirroring over an unlimited distance.
- Supports maximum sized logical volumes.
- Supports concurrent access configurations within the nodes on one site and allows you to have enhanced concurrent volume groups as geographically mirrored volume groups.

You can run concurrent instances of the same application on all nodes at the currently active site. That is, the data can be accessed concurrently within all the nodes of a single site with data being mirrored to the other site. However, data cannot be accessed concurrently at both sites at the same time.

Synchronous mirroring

Synchronous mirroring writes to both the local and remote sites at the same time, keeping your disaster recovery site up to date. While having both sites up to date is definitely advantageous, the time that it takes to write to the remote physical volumes can have a large impact on application response time.

There are two limiting factors when using synchronous mirroring: Network bandwidth and Network latency.

Network bandwidth

Network bandwidth is a limiting factor when the amount of data that needs to be sent over the network exceeds the network's capacity. If the network (or networks if there is more than one) is at full capacity, network buffers and queues fill up and messages must wait to be sent. When this happens, reads and writes to remote physical volumes take even longer and application response time suffers. While this might be acceptable for brief periods of peak activity when running batch applications or non-critical interactive applications, it is not

acceptable for most mission critical applications. The application is perceived as hanging, when in fact it is just waiting for remote physical volume I/O to complete.

A network bandwidth problem can be resolved by upgrading to a higher bandwidth network or by adding an additional data mirroring network - the RPV device driver supports up to four data mirroring networks. However, with synchronous mirroring, you need to provide enough network bandwidth to handle the data mirroring workload at its peak, in order to ensure acceptable response time. If these peaks do not occur very often and if there is a big difference between how much bandwidth is needed to handle the peaks versus how much bandwidth is required most of the time, then you may end up paying for a more expensive network without realizing its full benefits. For example, having a network that is only performing at maximum capacity for 10 percent of the time, just so that it can handle one hour of peak activity each day.

Network latency

Network latency is the time that it takes for messages to go across the network. Even when there is plenty of network bandwidth, it still takes some amount of time for the bits to travel over the communication link. The speed of the network is limited by the laws of physics; the greater the distance between the sites, the higher the amount of network latency. Even if a network is capable of transmitting data at a rate of 100 miles per millisecond, that still adds up over a long distance. For example, if the sites are 100 miles apart, a remote physical volume I/O request must travel 100 miles from the RPV client to the RPV server. After the disk gets updated, the result of the I/O request must travel 100 miles from the RPV server back to the RPV client. This 200 mile round trip adds 2 milliseconds to each remote physical volume I/O request, and this time can be much greater if there are routers or gateways on the network. Now suppose that the sites are 3000 miles apart. Each I/O request requires a 6000 mile round trip on the network, which adds approximately 60 milliseconds to each I/O request. This results in poor application response time that would likely result from having each logical volume write request take such a long time to complete. Synchronous mirroring is practical for a maximum of roughly 100 or 200 miles. Anything beyond that usually requires some form of asynchronous mirroring.

Asynchronous mirroring

Asynchronous mirroring allows the local site to be updated immediately and the remote site to be updated as bandwidth allows. The information is cached and sent later, as network resources become available. While this can greatly increase application response time, there is some risk of data loss.

Network bandwidth

When synchronous mirroring is used, you need to provide enough network bandwidth to handle the data mirroring workload at its peak, in order to ensure acceptable response time. However, when asynchronous mirroring is used, you may only need to provide enough network bandwidth for slightly more than the average amount of the data mirroring workload. It really depends on how much the peak differs from the average, and whether the production site cache is large enough to hold the excess write requests during the peak periods. In most cases, asynchronous mirroring requires a less expensive, lesser bandwidth network than synchronous mirroring. For example, if a synchronous solution requires a network that is only 10 percent utilized most of the time, but the same workload can be mirrored asynchronously over a low bandwidth network that is 75 percent utilized most of the time, then asynchronous mirroring may be a better choice than synchronous mirroring.

Network latency

Asynchronous mirroring allows data mirroring at the disaster recovery site to lag behind application writes that occur at the production site. This can greatly improve application response time by having AIX® LVM tell the application that its write has completed after the data has been written to the local disks, but without having to wait for the data to be written to the remote disks. The remote physical volume write requests are cached at the production site and mirrored to the disaster recovery site over a longer period of time, effectively removing the effects of network latency, which in turn allows the sites to be much farther apart without impacting application response time.

If the remote data mirroring is able to keep up with demand enough to prevent the cache from filling up, then there may not be a noticeable delay in application response time. However, once the caching limit is reached, application writes will have to wait until there is space in the cache. In a write-intensive application workload, the remote mirroring would quickly reach the cache limit and application response time would decrease. In such an environment, asynchronous mirroring does not offer any improvement over synchronous mirroring and, because of the risk of data loss, is not the best choice for mirroring.

Preventing data loss

Asynchronous mirroring creates the possibility of some amount of data loss from a production site disaster. If the remote site mirroring lags behind the local site, then you run the risk of losing that cached data in the event of a disaster. You need to determine how much data you are willing to risk losing.

Remote physical volume write requests are cached in permanent storage at the production site until they are written to disk at the disaster recovery site. After a node crash, you can recover these write requests. For example, suppose that a node crashes while it has the volume group varied online. You can recover the crashed node, bring the volume group back online, and have the asynchronous mirroring pick up where it stopped, with no more data loss than when using ordinary volume groups.

If you stop the application workload and take a volume group offline, all outstanding remote physical volume writes are written to disk at the remote site. For example, if you take the production site down for planned maintenance, you do not want the volume group to be brought online at the disaster recovery site while there are still outstanding writes sitting in the cache at the production site. By forcing the remote site to be brought up to date at the time that the volume group is taken offline, the application workload avoids accessing back-level data by mistake. Additionally, graceful PowerHA® SystemMirror® failover of asynchronously geographically mirrored volume groups from the production site to the disaster recovery site can take place without any data loss. The drawback to this approach is that it will take longer for the volume group to be taken offline, when the cache contains a backlog of remote physical volume write requests. Depending on how big the backlog is, it can take a very long time for all of the writes in the cache to be written to the disks at the remote site. And this, in turn, can cause all types of graceful failovers, whether they are local peer or site failovers, to take a very long time.

The only time when data loss may occur while using asynchronous mirroring, beyond what would be expected when using ordinary volume groups, is when the entire production site suddenly fails, before the mirroring to the disaster recovery site has had a chance to catch up. Whether or not the data is really lost depends on the circumstances of the failure, and in some cases how you want to deal with those circumstances. For example, a flood or fire can destroy all of the hardware at the production site. In that scenario, data loss would almost certainly occur. The lost data would consist of all the non-mirrored remote physical volume writes that were in the cache on the production site at the time of the failure. In another situation, a power outage can bring down the entire production site without destroying any hardware. In this scenario, the data is still there, but it cannot be accessed until the power can be restored and the system can be brought back online. You can choose to wait for the production site to be recovered, so you can avoid losing the non-mirror data or you can move your application workload to the disaster recovery site, with some amount of data loss.

Data divergence

Data divergence is a state where each site's disks contain data updates that have not been mirrored to the other site. For example, if a disaster destroys the disks at the production site, then the only copy of the data exists at the disaster recovery site. Using asynchronous mirroring, the possibility exists that the data will be back level, due to data caching. However, it is possible for the production site to fail without hardware damage. In this case, the data is still there, but cannot be accessed until the production site can be brought back online. In this case, you can wait for the production site to be brought back online or you can move the application workload over to the disaster recovery site. If you move the application workload over to the disaster recovery site, you risk data divergence as the application begins to use the back-level data on the disaster recovery site disks. You will need to determine what action PowerHA® SystemMirror® should take should the production site go down while the disaster recovery site contains back-level data.

Once data divergence occurs, you will have to decide how you want to recover. If very few or no transactions are made at the disaster recovery site, you should be able to move back to production site with few complications. However, if your disaster recovery site has been running your applications for a long time, you cannot simply go back to your production site without risking some of your data.

GLVM for a PowerHA® SystemMirror® cluster

With GLVM for PowerHA® SystemMirror® Enterprise Edition you can include geographically mirrored volume groups in a PowerHA® SystemMirror® cluster.

This function requires two PowerHA® SystemMirror® sites and a PowerHA® SystemMirror® component to which you can assign nodes. Cluster nodes access the same geographically mirrored volume groups, but the nodes at each site access them from local and remote physical volumes. This process is different from a single-site PowerHA® SystemMirror® environment, in which all cluster nodes share volume groups that have physical connections to the same disks.

PowerHA® SystemMirror® sites

PowerHA® SystemMirror® supports two sites. The local site is the active site, and the remote site is the backup site.

The terms *local site* and *remote site* are relative. The site on which the resource group is hosted under normal cluster conditions is the local site. In other words, the home node for this resource group belongs to the local site. The site to which the resource group can be moved during a failover (if no other nodes at the local site can host the resource group) is the remote site for this resource group.

The Inter-Site Management Policy for a resource group directs how a resource group and its resources fall over in response to a site outage, and how they fall back if configured to do so. For each resource group, one site is an active production site and the other a backup site. If the nodes at the active production site become unavailable, the backup site becomes the active production site.

Each site contains at least one node.

The resource groups that contain geographically mirrored volume groups have two types of management policies:

- Resource group management policies determine failover behavior if a node becomes unavailable. The resource group policies for GLVM for PowerHA® SystemMirror® Enterprise Edition can be concurrent or nonconcurrent.
- Site management policies determine failover and fallback behavior if all of the nodes at a site are not available. GLVM for PowerHA® SystemMirror® Enterprise Edition site management policies must be nonconcurrent. These policies keep the resource groups with geographically mirrored volume groups highly available.

Fallover and fallback

GLVM for PowerHA® SystemMirror® Enterprise Edition handles the automation of failover from one site to another in response to an outage at a local site, minimizing recovery time for the application. When a site fails, the resource group configuration ensures that the geographically mirrored volume groups are accessible from the remote site.

GLVM for PowerHA® SystemMirror® Enterprise Edition automates application recovery by managing:

- The failover of nodes within a site based on node priority (as identified in the node list for a resource group)
- The failover between sites (as specified by the site management policy for a resource group)
- The fallback of a resource group to a node as configured.

When an application is running on a local (or primary) site and the remote site is part of the cluster:

- Updates to the application data are made to the RLVs associated with the geographically mirrored volume groups
- Data is mirrored using the GLVM on the RLVs.

If the node or the disks at the local site become unavailable:

- The application moves to a server at the remote site (if no other nodes at the local site can host the application).
- The application continues operation using the mirrored copy of the data.

When the initial local site becomes active again, resource group and site management policies determine whether or not the application moves back to the local site:

- The direction of mirroring may be reversed.
- The application may be stopped and restarted on another node.

Note: The initial local site may become the new remote site with the direction of mirroring reversed.

Sample configuration

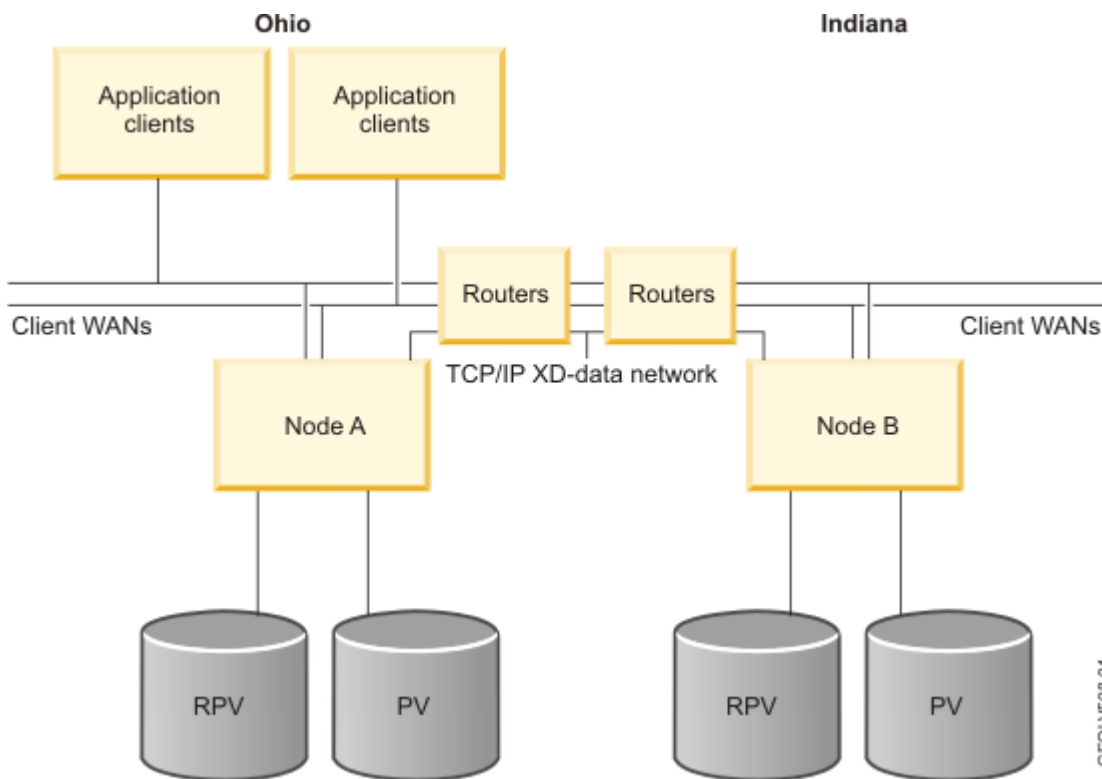
GLVM for a PowerHA® SystemMirror® cluster consists of two sites located at a geographic distance from one another. Both sites can support a combined total of 16 cluster nodes.

There are two types of disaster protection:

- **Fallover to a Remote Backup Site.** In this configuration, a local site serves as a production site and a remote site serves as a backup site. The backup site includes hardware, system and application software, and application data and files. It is ready to take over the current workload on the local site. In the event of a failure, the application from the local site automatically falls over to the remote site.
- **Remote Mutual Takeover.** Remote mutual takeover allows two sites to serve as backups for each other. If a failure occurs at a node on one site, the node at the other site automatically takes over the application. With mutual takeover, the initial local site may become the new remote site with the direction of mirroring reversed. Mutual takeover requires two separate resource groups.

Cluster configuration with sites and RPVs

The following figure shows a two-site cluster configuration with one node at each site:



In this example:

- **Sites and nodes configuration.**
 - The PowerHA® SystemMirror® cluster contains two sites: site Ohio and site Indiana.
 - Each site has one node.
 - Each node has a local physical volume configured.
- **Networks configuration.**
 - The TCP/IP XD_data network is used for data mirroring. This diagram uses just one network. You can configure up to four XD_data mirroring networks in a cluster with GLVM for PowerHA® SystemMirror® Enterprise Edition.
- **RPV configuration.** Two RPV client and server pairs are configured for each resource group. Note, this is true in this example because the resource group contains two disks. In general, however, an RPV client and server pair must be configured for each disk. The node on which the disk is local requires an RPV server, and the node on which the disk remote requires an RPV client.

- **RPV client on Node A and RPV server on Node B.** For a resource group residing on site Ohio, site Ohio is the local site and site Indiana is the remote site. Node A has an RPV client configured and Node B has a corresponding RPV server configured. The mirroring occurs in the direction from the RPV client on the local site to the RPV server on the remote site. This enables data mirroring for volume groups and applications residing on site Ohio (the local site for this resource group) to site Indiana (the remote site for this resource group).
- **RPV client on Node B and RPV server on Node A.** For a resource group residing on site Indiana, site Indiana is the local site, and site Ohio is the remote site. Node B has an RPV client configured. Node A has a corresponding RPV server configured. For this resource group, the applications that reside on site Indiana are mirrored to site Ohio. The mirroring occurs in the direction from the RPV client configured on Node B at site Indiana to the RPV server configured on Node A at site Ohio.

RPV client-server pairs required at both sites

The previous example illustrates that configuring RPV client-server pairs at both sites is necessary to ensure mutual takeover and recovery in PowerHA® SystemMirror® for resource groups that reside on nodes at each site.

For instance, a resource group resides on the local site and Node A is its home node. The goal of GLVM for PowerHA® SystemMirror® Enterprise Edition is to recover the resource group on the remote site, in case of a failure on the local site, and if there is no other node at the local site on which to recover the resource group.

For this purpose, an RPV client-server pair must be configured for the volume group in the resource group. An RPV client is configured on Node A (local site for this resource group) and the corresponding RPV server is configured on Node B (remote site for this resource group). Geographic mirroring occurs in the direction from the RPV client on the local site to the RPV server on the remote site. If a failover between sites occurs, GLVM for PowerHA® SystemMirror® Enterprise Edition recovers the resource group on Node B and the application continues to run. GLVM for PowerHA® SystemMirror® Enterprise Edition activates the local volume group and enables access to the disk directly on the takeover node.

Similarly, for another resource group that initially resides on site Indiana (this is the local site for this resource group) and has Node B as its home node, an RPV client must be configured on Node B and the corresponding RPV server must be configured on Node A. If a failover between sites occurs, GLVM for PowerHA® SystemMirror® Enterprise Edition recovers the resource group on Node A and the application continues to run. GLVM for PowerHA® SystemMirror® Enterprise Edition activates the local volume group and enables access to the disk directly on the takeover node at another site.

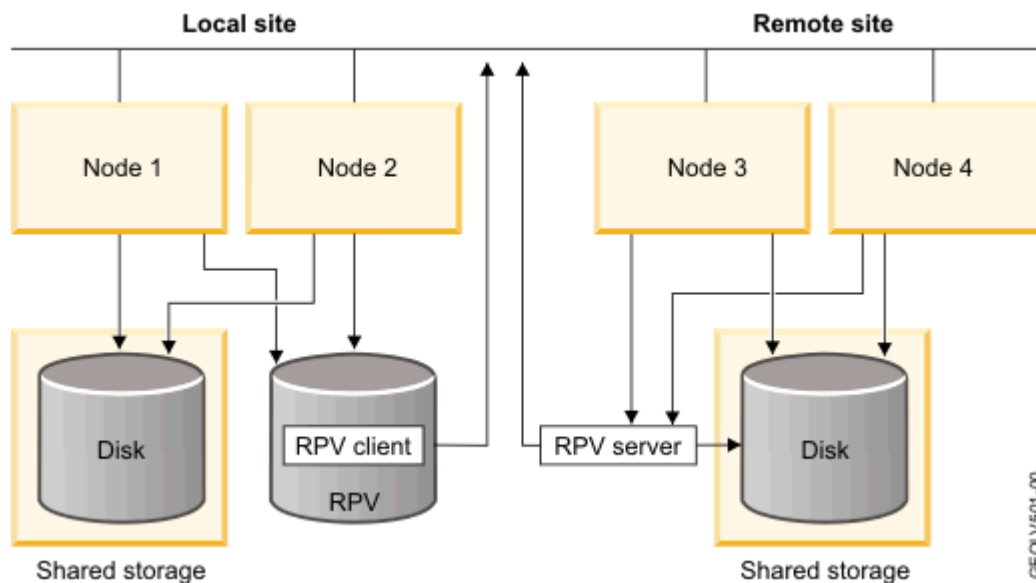
Cluster configuration with two nodes at each site

The following figure shows a sample two-site two-node cluster configuration with nodes that use geographically mirrored volume groups as shared disks.

In this example:

- The PowerHA® SystemMirror® cluster contains two sites: the local site and the remote site.
- Each site has two nodes.
- Node 2 on the local site has an RPV client configured. (This node also must have an RPV server configured, but it is not shown. See the following note).
- Node 3 on the remote site has an RPV server configured that communicates with the RPV client on Node 2 using at least one IP-based network (*XD_data*). A *XD_ip* networks is not shown. You can have up to four *XD_data* networks. The RPV client on Node 3 is also not shown, see the following note.

Note: In this example, only one pair of RPV server and client is displayed. However, to ensure that GLVM for PowerHA® SystemMirror® Enterprise Edition provides high availability by recovering the resource group, Node 3 must also have another RPV client configured (not shown), and Node2 must have a corresponding RPV server configured (not shown).



In this configuration, a resource group on Node 2 contains an application that accesses volume groups at the local site. All volume groups that are accessed by this application have both local PVs and RPVs. These volume groups are geographically mirrored with the GLVM mirroring function. The groups use an RPV as a layer that allows the node to access them.

Example: From ordinary LVM mirroring to GLVM

This example considers a typical case scenario with an application that contains mission-critical data. The initial configuration contains several data storage disks at one geographic location. The case description walks you through adding additional remotely located disks and servers to the configuration and explains how GLVM for PowerHA® SystemMirror® Enterprise Edition ensures data availability and disaster recovery.

In this example, the new terms, *Geographic Logical Volume Manager (GLVM)* and *remote physical volume (RPV)* are used and defined. Also, this example displays how PowerHA® SystemMirror® recovery features fit into High Availability.

Consider a single system that is running a mission-critical airline reservation system. All of the data resides on two disks, PV1 and PV2. Each disk contains exactly half of the data. In other words, there is only one copy of the data which spans both disks.

At this stage, the airline reservation system configuration has many single points of failure. The failure of either disk will cause an outage, until the failed disk can be replaced and the latest backup data can be restored. This problem can be avoided by using the AIX® LVM to create a second (mirror) copy of the data.

Let us assume that two more disks are added to this system. Now the original data resides on disks PV1 and PV2, and another copy of the data resides on disks PV3 and PV4. When the AIX® LVM mirroring is used, a disk is no longer a single point of failure. However, there are still other single points of failure that need to be addressed. The airline reservation system still has many hardware and software components which may fail and cause an outage. Adding a PowerHA® SystemMirror® cluster is the best way to eliminate most single points of failure.

Let us assume that a second system is added, and the two servers are now combined into a PowerHA® SystemMirror® cluster. Each server is physically connected to all four disks. If the server that is running the airline reservation system fails, PowerHA® SystemMirror® detects this and recovers the application by varying on the mission-critical volume group and starting the applications on the surviving node.

For clarity, let us label the servers (nodes) and disks as Node1 and Node2. Since disk names are not usually the same across the entire cluster, each node may know the disks by different names. The configuration is now as follows:

| Disks | Node 1 | Node 2 |
|-------|---------|---------|
| PV1 | hdisk11 | hdisk12 |
| PV2 | hdisk12 | hdisk13 |
| PV3 | hdisk13 | hdisk14 |

| Disks | Node 1 | Node 2 |
|-------|---------|---------|
| PV4 | hdisk14 | hdisk15 |

Now assume that Node 2, along with disks PV3 and PV4 is moved several hundred miles away. The two servers are now far away from each other, connected by an ordinary TCP/IP-based Wide Area Network. A PowerHA® SystemMirror® cluster can exist with its nodes at two geographically separated sites. The biggest challenge is that most disk subsystems do not allow a server to be physically attached to a disk that is located several hundred miles away.

The Geographical Logical Volume Manager solves this problem by mirroring data over the Wide Area Network. GLVM uses the RPV device driver, which is a pseudo device driver that allows the AIX® LVM to access remote disks as if they were locally attached.

The following description explains the hardware configuration in detail and shows where the RPV pseudo devices fit in the configuration. Suppose that Node 1 resides in Poughkeepsie, NY, and Node 2 resides in Philadelphia, PA:

| Poughkeepsie, NY | | Philadelphia, PA | |
|------------------|---------|------------------|---------|
| Disks | Node 1 | Disks | Node 2 |
| PV1 | hdisk11 | PV3 | hdisk14 |
| PV2 | hdisk12 | PV4 | hdisk15 |

First, we add RPV servers to node2 for disks PV3 and PV4.

An *RPV server* is a software entity which processes disk I/O requests from a remote RPV client. For example, rpvserver0 will process disk I/O requests from Node 1:

| Philadelphia, PA | |
|------------------|--------------------|
| Disks | Node 2 |
| PV3 | hdisk14 rpvserver0 |
| PV4 | hdisk15 rpvserver1 |

Next, we add RPV clients to Node 1 for disks PV3 and PV4:

| Poughkeepsie, NY | |
|------------------|-----------------------------------------|
| Disks | Node 1 |
| PV1 | hdisk11 |
| PV2 | hdisk12 |
| PV3 | hdisk13 (Remote Physical Volume Client) |
| PV4 | hdisk14 (Remote Physical Volume Client) |

The configuration on Node 1 now includes two additional disk devices that are known as hdisk13 and hdisk14. The main difference between these disk devices and local physical volumes is that these disk devices are RPV clients.

An *RPV client* is a software entity which appears to the system as a disk device driver. It redirects disk I/O requests to a remote RPV server, instead of a locally attached disk. When the LVM opens hdisk13 and reads and writes to it, all of the disk I/O requests are redirected to rpvserver0 on Node 2 located in Philadelphia. The RPV server processes each I/O request and sends its result back to the RPV client, which in turn passes it back to the LVM.

In this configuration, you can locally (at the Poughkeepsie site) vary on a volume group which has one mirror copy in Poughkeepsie and a second mirror copy in Philadelphia. Now, you can add RPV servers to Node 1 for disks PV1 and PV2:

| Poughkeepsie, NY | |
|------------------|-----------------------------------------|
| Disks | Node1 |
| PV1 | hdisk11 rpvserver5 |
| PV2 | hdisk12 rpvserver6 |
| PV3 | hdisk13 (Remote Physical Volume Client) |

| | |
|-------------------------|-----------------------------------------|
| Poughkeepsie, NY | |
| PV4 | hdisk14 (Remote Physical Volume Client) |

You also add RPV clients to Node 2 for disks PV1 and PV2:

| | |
|-------------------------|-----------------------------------------|
| Philadelphia, PA | |
| Disks | Node2 |
| PV1 | hdisk12 (Remote Physical Volume Client) |
| PV2 | hdisk13 (Remote Physical Volume Client) |
| PV3 | hdisk14 rpvserver0 |
| PV4 | hdisk15 rpvserver1 |

Creating this second set of remote physical volumes allows the volume group to be varied on in Philadelphia, with data mirroring occurring in the opposite direction. (Of course, the volume group cannot be varied on in Poughkeepsie at the same time).

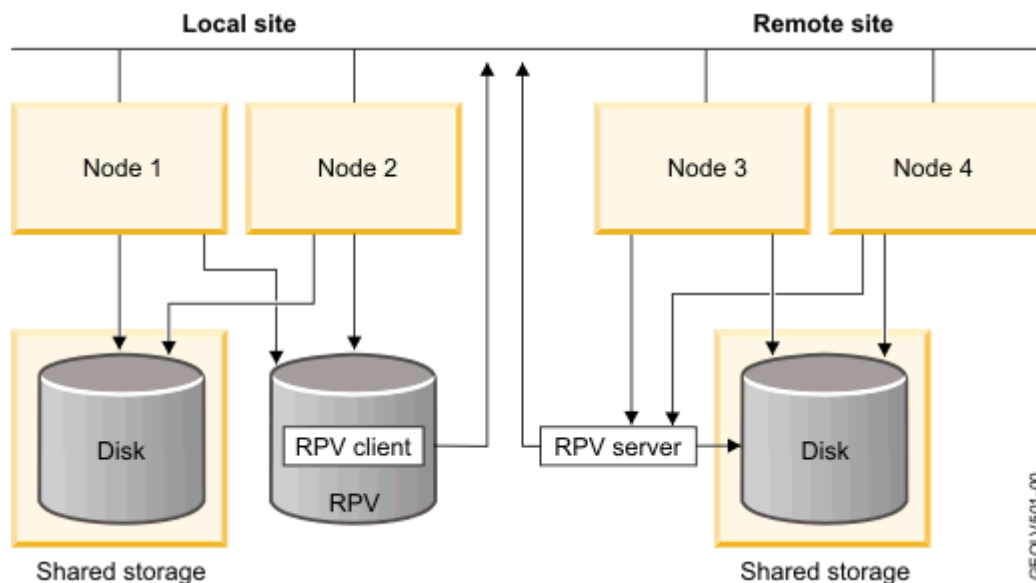
Once each node is able to access all four disks, it becomes possible to create a volume group which is mirrored across both sites. However, you have to be careful to make sure that each mirror copy is contained within a single site. You would not want to store the first mirror copy on disks PV1 and PV3, and then store the second mirror copy on disks PV2 and PV4. You could end up with two copies of the same data at each site. In this case, after a site disaster occurs, you would end up with an incomplete copy of the data on the surviving site. Fortunately, GLVM for PowerHA® SystemMirror® Enterprise Edition provides the GLVM Utilities to help you to create the geographically mirrored volume groups and logical volumes correctly.

To summarize, combining the RPV device driver with the AIX® LVM allows volume groups and logical volumes to span two sites. A site failover in response to a disaster is conceptually similar to a node failover in a traditional PowerHA® SystemMirror® cluster.

Example: Node and site failure

This topic describes how GLVM for PowerHA® SystemMirror® Enterprise Edition handles failures and what happens in the cluster to ensure that the geographically mirrored data remains available to the application.

The following figure with two nodes at each site is used in these scenarios. In this configuration, Node 2 is the home node, or the primary owner of the resource group. It has the RPV client configured and resides on the local site. Node 1 is the second priority node for the resource group:



Node failure on the local site

If Node 2 fails, the resource group with the application falls over to Node 1.

At a high level, on Node 1, GLVM for PowerHA® SystemMirror® Enterprise Edition detects that Node 2, the primary owner of the resource group, has failed and moves the resource group to Node 1. It also activates the geographically mirrored volume group on Node1 and ensures that the RPV client on Node 1 is available for communication with the RPV server on Node 3.

The application is kept highly available and the data continues to be geographically mirrored. The application continues to send I/O requests to the RPV, and the GLVM function ensures that the RPV client and RPV server continue to communicate.

If Node 2 rejoins the cluster, based on the resource group policy PowerHA® SystemMirror® performs the resource group fallback. The resource group moves back to Node 2 (depending on the resource group fallback policy) and the RPV client becomes available on Node 2. This enables reestablishing communication between the RPV client and server. The resource group resumes mirroring data through the RPV.

Node failure on the remote site

If Node 3 that has the RPV server configured, fails, GLVM for PowerHA® SystemMirror® Enterprise Edition moves the backup instance of the resource group to Node 4 on the remote site. PowerHA® SystemMirror® also ensures that the RPV server is now available on Node4.

The application is kept highly available and the data continues to be geographically mirrored. The application continues to send I/O requests to the RPV, and the GLVM function ensures that the RPV client and RPV server continue to communicate.

When the node at the remote site rejoins the cluster, based on the resource group policy PowerHA® SystemMirror® performs the resource group fallback.

Local site failure


If, as a result of a failure, there is no node available on the local site to host a resource group with the geographically mirrored volume group and the application, the resource group falls over to one of the nodes on the remote site.

For example, the resource group falls over from Node 2 on the local site to Node 3 on the remote site. In this case, GLVM for PowerHA® SystemMirror® Enterprise Edition forcefully activates (if configured to do so) the local volume group with the mirror copy of the data on Node 3, and Node 3 acquires the resource group. The application in the resource group accesses that data directly and does not use an RPV.

If Node 2 rejoins the cluster, PowerHA® SystemMirror® performs the resource group fallback (depending on the resource group fallback policy). The resource group moves back to Node 2 and the RPV client becomes available on Node 2. The communication between the RPV client and server is reestablished. GLVM for PowerHA® SystemMirror® Enterprise Edition synchronizes the mirror copies and reestablishes data mirroring. The resource group resumes mirroring data through the RPV.

Remote site failure

The following figure shows a remote site failure (nodes at each site are not shown):



Note: The previous figure shows the cluster networks.

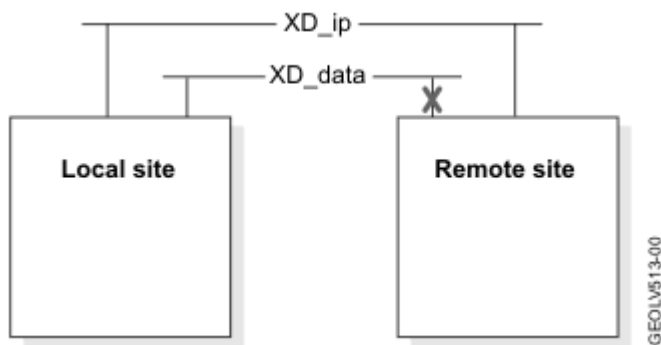
If, as a result of a failure, there is no node available on the remote site to host the RPV server for the resource group with the geographically mirrored volume group, PowerHA® SystemMirror® stops the RPV client from trying to communicate with the RPV server.

If a node that hosted the RPV server on the remote site rejoins the cluster, GLVM for PowerHA® SystemMirror® Enterprise Edition resumes the RPV server on that node and reestablishes the communication between the RPV client and server. Mirroring is reestablished too, and the mirror copies are synchronized. The resource group remains on Node 2 and the RPV client begins to communicate with the RPV server on Node 3. The resource group resumes mirroring data through the RPV.

Preventing cluster partitioning

To prevent cluster partitioning, configure a network for heartbeating between the sites, in addition to several (up to four) IP-based networks over which the mirroring data is being transferred. If all mirroring network connections between the RPV client and server (and between the sites) fail, the heartbeating network prevents cluster partitioning and the resultant data divergence.

The following figure illustrates the failure of the data mirroring network in the cluster with two sites (nodes at each site are not shown):



Planning for GLVM

Before you can start using GLVM for PowerHA® SystemMirror® Enterprise Edition you must plan on how you will implement it into your environment.

Planning geographically mirrored volume groups

To plan for using geographically mirrored volume groups in a PowerHA® SystemMirror® cluster, you must complete the planning tasks for GLVM and planning tasks for PowerHA® SystemMirror® integration with GLVM.

You should be familiar with the planning tasks for PowerHA® SystemMirror®.

Review the following planning information before you consider implementing GLVM for PowerHA® SystemMirror® Enterprise Edition:

- You can have several IBM® servers with AIX® installed, but no PowerHA® SystemMirror® cluster configured. In this case, you first plan geographically mirrored volume groups and install GLVM for PowerHA® SystemMirror® Enterprise Edition. After configuring geographically mirrored volume groups, you can include them into a PowerHA® SystemMirror® cluster on two sites.
- Alternatively, you can have a PowerHA® SystemMirror® cluster already configured, and wish to implement a disaster recovery solution between two sites that uses the mirroring function provided by GLVM for PowerHA® SystemMirror® Enterprise Edition. In this case, you should plan for the geographically mirrored volume groups while keeping in mind your already configured cluster resources, networks and sites. You might need to add a second site or nodes to your cluster. You also might need to extend the existing volume groups to it.

If you plan on turning your existing enhanced concurrent mode volume groups into geographically mirrored volume groups, verify that PowerHA® SystemMirror® is running on the nodes (to enable you to perform some operations that won't run in AIX® without PowerHA® SystemMirror® running on the nodes).

After planning and configuring the geographically mirrored volume groups, you can add them to the resource groups in the PowerHA® SystemMirror® cluster.

Note: You might need to create new resource groups if one is not available.

Requirements and limitations

GLVM for PowerHA® SystemMirror® Enterprise Edition has several different requirements and limitations.

The follow are requirements and limitations for geographically mirrored volume groups:

- You must set the inter-disk allocation policy for logical volumes in AIX® to Super Strict. The Super Strict policy enforces that a complete mirror copy resides on separate sets of either local or remote physical volumes. Since the set of RPVs exists on one remote site, the Super Strict policy for mirroring and GLVM for PowerHA® SystemMirror® Enterprise Edition ensure that when you use the GLVM function to create geographically mirrored volume groups, you always have a complete copy on the remote site.
- You can only plan and configure up to three copies of the logical volumes. You can configure at least one mirror copy of the logical volume at each site. One of the sites may optionally contain a second copy, that is, have a total of up to two copies.

Note: If you do not plan to configure mirror copies between geographically remote sites, all three copies can reside on the nodes at the same geographic location.

- You can only plan and configure two sites, one local and one remote site. Site names must correspond with the site names in PowerHA® SystemMirror®. All nodes configured at each site must be members of the site.
- The **rootvg** volume group cannot be geographically mirrored.
- You must use nonconcurrent or enhanced concurrent mode volume groups. For enhanced concurrent volume groups that you also want to make geographically mirrored, ensure that PowerHA® SystemMirror® cluster services are running before adding and managing RPVs.
- You must use either a nonconcurrent or concurrent resource group management policy and a nonconcurrent inter-site management policy for resource groups that contain geographically mirrored volume groups. You cannot create dependent resource groups.
- You can only configure up to four IP-based networks, XD_data, in PowerHA® SystemMirror® that will provide the data stream between mirrored volume group copies at two sites. In addition, to prevent cluster partitioning, you can also configure a XD_ip network between sites.

You can use the IP Security (IPsec) feature to secure the RPV client-server network traffic between the sites over an XD_data network.

Software Requirements

GLVM for PowerHA® SystemMirror® Enterprise Edition requires specific versions of AIX® and RSCT. The RSCT filesets are shipped with the PowerHA® SystemMirror® installation media and are installed automatically. The PowerHA® SystemMirror® Enterprise Edition software uses 1 MB of disk space. Ensure that the **/usr** file system has 1 MB free disk space.

Planning for GLVM for PowerHA® SystemMirror® Enterprise Edition

There are several steps you should perform when planning to use GLVM for PowerHA® SystemMirror® Enterprise Edition.

To correctly plan to use GLVM for PowerHA® SystemMirror® Enterprise Edition, complete the following steps:

1. Identify the remote physical volume (RPV) sites and nodes that belong to each site.
2. Identify the volume groups that you plan to configure as geographically mirrored volume groups.
3. Plan the RPV clients and servers for each physical volume in a volume group.
4. Plan connections for the GLVM mirroring function between sites and for the site connections to use security functions.

Planning for sites

A *site name* is a name attribute that you assign to the RPV server. The RPV device driver uses a site name attribute to identify a set of nodes that reside at the same physical location.

To successfully plan for sites, review the following information.

- You have two geographically-distant sites: one local site, and one remote site. Each site obtains a unique site name.
- All nodes that have access to the same physical volumes must be part of the same site.
- RPV site names and PowerHA® SystemMirror® site names must be the same.

Planning geographically mirrored volume groups

In general, the AIX® LVM allows several ways to create local mirrors of the volume groups.

To create geographically mirrored volume groups, review the following information.

- Ensure you have sufficient disk space on the nodes at each site to mirror your physical volumes.
- Plan to create the volume groups. In PowerHA® SystemMirror® Enterprise Edition for GLVM, in addition to standard volume groups, you can have enhanced concurrent volume groups and turn them into geographically mirrored volume groups by adding RPVs. Note that to add enhanced concurrent volume groups that are also geographically mirrored, you need to start with an already configured PowerHA® SystemMirror® cluster, and have the cluster services running on the nodes at both sites.
- Plan to create the logical volumes with the Super Strict inter-disk allocation policy. These logical volumes must be associated with local physical volumes and belong to the same site.

Planning remote physical volumes

After you create a logical volume in AIX®, you can use the GLVM Utilities SMIT interface to add mirror copies. To allow for geographic mirroring, the logical volume requires a mirror copy that will reside on the nodes at the local site and a mirror copy that will reside on the nodes at the remote site.

The RPVs are viewed by the LVM as regular physical volumes. In particular, when answering to a write request, there is no difference to the system whether a logical volume copy resides on the local or remote physical volume.

RPV client-server pairs are configured for the physical volumes which comprise a volume group that contains one or more logical volumes that are mirrored across two sites.

Note that each logical volume does not require its own RPV. Multiple logical volumes may share an RPV. You do not need to have a one-to-one correspondence between the number of local physical volumes and remote physical volumes, although you must have an adequate number of RPVs to satisfy the requirements for mirroring all logical volumes.

Each RPV client communicates with its corresponding RPV server over one or more IP-based XD_data networks. To ensure load balancing, an RPV client sends I/O requests across alternating networks: An RPV client makes only one attempt to access its corresponding RPV server over the same network, and if this attempt is not successful, it proceeds to try the connection over the next available network.

Similarly, each RPV server communicates with its corresponding RPV client over one or more IP-based data mirroring networks.

For instance, consider the following scenario with four 1GB disks at site A, and one 4GB disk at site B. In this case, the nodes are not mentioned for simplicity.

In this case, at site B, you plan to have one RPV server configured for the 4GB disk, and at site A, the corresponding RPV client. In addition, at site A, you plan for the configuration of four RPV servers for each of the 1GB disks, and, at site B, for the configuration of the corresponding four RPV clients. You also plan to have more than one data mirroring network.

Planning networks

You need to plan for a few different cluster networks for GLVM for PowerHA® SystemMirror® Enterprise Edition.

The following list displays the cluster networks for GLVM for PowerHA® SystemMirror® Enterprise Edition.

- IP-based networks
- Back up connection by using an IP link between sites

Determining the number of data mirroring networks you need

You should check with your network provider the characteristics of the IP network that is provisioned to you, and based on this decide whether you need to have one or more XD_data networks. The goal is to eliminate a single point of failure. For instance, if you know that multiple physical networks are not implemented for the XD_data network that is provisioned to you, plan to obtain and configure more than one XD_data networks in GLVM for PowerHA® SystemMirror® Enterprise Edition.

Alternatively, if you know that the network provider implements your network with multiple physical networks underneath, this already ensures that there are no single points of failure for the physical network that supports the XD_data network, and in this case, it is sufficient to have just one XD_data network.

In the cluster configuration, GLVM for PowerHA® SystemMirror® Enterprise Edition lets you have up to four IP-based networks for RPV client-server communication and for heartbeating. For these networks, plan to allocate service IP addresses and boot IP addresses. Also, if you are planning to have more than one XD_data network, ensure that they have similar bandwidth and latency to enable load balancing between them.

Planning the networks that will not be used for data mirroring

Plan to have at least one additional IP-based network for heartbeating to prevent cluster partitioning and subsequent data corruption.

If you are planning to configure an additional IP-based network that won't be used for data mirroring, plan to allocate service and back up IP addresses for it.

Planning for network security

You can use the IP Security (IPsec) feature of AIX® to ensure additional security on the networks over which the traffic between the RPV servers and clients is sent.

The IP Security feature of AIX® is installed separately from the other AIX® functionality. Use the **System Management > Communications Applications and Services > TCP/IP > Configure IP Security (IPv4)** option in SMIT to configure IP security. You can also use the `smit ipsec4 fastpath`.

You can also configure VPN tunnels between the sites that use PowerHA® SystemMirror® Enterprise Edition for GLVM mirroring.

Planning for integrating geographically mirrored volume groups into the cluster

When planning for integrating geographically mirrored volume groups into the cluster, there are several tasks that you need to perform.

To successfully plan for integrating geographically mirrored volume groups into the cluster, complete the following high-level tasks.

1. Plan the PowerHA® SystemMirror® sites. The RPV servers site names and PowerHA® SystemMirror® site names must match.
2. Plan the PowerHA® SystemMirror® networks. You will define the networks used by GLVM for PowerHA® SystemMirror® Enterprise Edition to the PowerHA® SystemMirror® cluster configuration.
3. Identify the volume groups to be included in specified PowerHA® SystemMirror® resource groups.
4. Plan which resource groups will contain the geographically mirrored volume groups.

Planning PowerHA® SystemMirror® clusters with geographically mirrored volume groups

The following topics provide planning information for PowerHA® SystemMirror® integration with GLVM.

Planning the PowerHA® SystemMirror® configuration

When you include the geographically mirrored volume groups into a PowerHA® SystemMirror® cluster, GLVM for PowerHA® SystemMirror® Enterprise Edition treats them as a geographically mirrored resource that will be kept highly available.

You need to plan for the following components in a PowerHA® SystemMirror® cluster that will include geographically mirrored volume groups.

- Two sites (only)
- Resource groups
- Networks.

Planning PowerHA® SystemMirror® sites

GLVM for PowerHA® SystemMirror® Enterprise Edition requires two PowerHA® SystemMirror® sites. Each PowerHA® SystemMirror® site must have the same name as the RPV server site name.

Planning for resource groups in an PowerHA® SystemMirror® cluster

In addition to basic resource group planning, you will need to plan for other resource group attributes.

The following list displays the attributes for resource groups in an PowerHA® SystemMirror® cluster.

- The node list for the resource groups must contain at least one node from one site (the local site for the resource group) and one node from another site (the remote site for the resource group).
- The resource groups cannot contain other, standard volume groups, in addition to the geographically mirrored volume groups.
- The resource group policy for a resource group that contains geographically mirrored volume groups can be either concurrent or nonconcurrent. The resource groups can contain both standard and enhanced concurrent mode volume groups. The inter-site management policy must be nonconcurrent, that is, it may not be set to Ignore or to Online on both Sites.
- You cannot configure RPV devices as raw physical volumes in PowerHA® SystemMirror® resource groups.
- You can use Dynamic Automatic Reconfiguration, or DARE, to dynamically change the resources in the resource groups containing the geographically mirrored volume groups. For information on what types of changes you can make, see the *Administration Guide*.
- GLVM for PowerHA® SystemMirror® Enterprise Edition processes the resource group with the geographically mirrored volume groups in parallel by default, unless you specify a serial processing order for them.

Planning for geographically mirrored volume groups in a PowerHA® SystemMirror® cluster

You add the geographically mirrored volume groups to the resource group using the **Volume Groups** field in the PowerHA® SystemMirror® **Change/Show Resources and Attributes for a Resource Group** SMIT panel.

- The field **Automatically Import Volume Groups** in the PowerHA® SystemMirror® **Change/Show Resources and Attributes for a Resource Group** SMIT panel must be set to **False** for geographically mirrored volume groups. If this option is set to **True** and the volume group contains RPVs, PowerHA® SystemMirror® issues a warning and does not automatically import the geographically mirrored volume group to the other nodes in the resource group.
- You can include a geographically mirrored volume group in a PowerHA® SystemMirror® resource group only if the volume group's associated RPV clients and servers are configured on nodes that belong to the cluster. If an RPV is configured on a machine outside of the PowerHA® SystemMirror® cluster, PowerHA® SystemMirror® does not allow you to include its associated volume group in the resource group.
- You must geographically mirror all of the logical volumes in a geographically mirrored volume group. In other words, you cannot choose to geographically mirror some logical volumes but not others in the same

volume group. All volumes that belong to the group must be mirrored remotely. However, a volume group may itself have a local copy, as well as a remote copy.

- In GLVM for PowerHA® SystemMirror® Enterprise Edition, you can configure a geographically mirrored volume group that is also an enhanced concurrent mode volume group. Using GLVM for PowerHA® SystemMirror® Enterprise Edition you can include RPVs into previously existing enhanced concurrent volume groups, however, to do so you must have PowerHA® SystemMirror® cluster services running on the nodes.
- Fast disk takeover and disk heartbeating are not supported on remote disks that are part of a geographically mirrored volume group.
- In GLVM for PowerHA® SystemMirror® Enterprise Edition, the C-SPOC utility does not allow managing the geographically mirrored volume groups from a single node. The GLVM Utilities SMIT interface does not provide a way to vary on, vary off, import or remove geographically mirrored volume groups from the cluster. (Note, it allows you to remove RPVs from a geographically mirrored volume group, thereby making it a regular volume group). Instead, you can vary on, vary off, import, change and remove geographically mirrored volume groups in the cluster on each node individually using the regular AIX® SMIT interface.
- Carefully consider quorum and forced varyon issues when planning the geographically mirrored volume groups.

Relinquishing control over RPVs in a PowerHA® SystemMirror® cluster

Once you configure resource groups that contain volume groups with RPVs, you relinquish control over RPVs to GLVM for PowerHA® SystemMirror® Enterprise Edition.

Once RPVs become part of volume groups that belong to an active PowerHA® SystemMirror® cluster, GLVM for PowerHA® SystemMirror® Enterprise Edition takes care of changing the RPV states (defined and available), during cluster events.

Note: If for some reason you stop the cluster services and must manage the RPVs manually, the RPVs must be configured in the defined state before you restart the cluster services. This enables management of RPVs by GLVM for PowerHA® SystemMirror® Enterprise Edition once the nodes join the cluster.

Site policy

When configuring a resource group that will include the geographically mirrored volume groups, the resource group includes nodes in both cluster sites. The resource group can contain both standard volume groups and enhanced concurrent mode volume groups that are geographically mirrored.

In order for GLVM for PowerHA® SystemMirror® Enterprise Edition to manage resources, you assign an inter-site management policy to each resource group. You can use Prefer Primary Site or Online on Either Site as the nonconcurrent inter-site policies for resource groups that include geographically mirrored volume groups.

Prefer Primary Site

In a two-site configuration, resource groups that contain geographically mirrored volume groups at startup reside on the site with the higher priority, fall over to the other site in case of a site failure, and then fall back to the site with the higher priority.

Online on Either Site

Resource groups with geographically mirrored volume groups reside on either site at startup, fall over to the other site in case of the site failure, and remain on that site after recovery.

Planning for PowerHA® SystemMirror® networks

You need to plan for the configuration of several different PowerHA® SystemMirror® networks.

The following list displays the PowerHA® SystemMirror® networks you need to plan on configuring.

- Up to four IP-based PowerHA® SystemMirror® networks that are used for mirroring data between all the nodes at both sites on which geographically mirrored volume groups reside. These networks have a type of XD_data.
- Also plan for having a XD_ip network. Having this network in the cluster helps to prevent cluster partitioning.

Note: NFS cross-mounting is not supported on XD networks.

Networks overview

GLVM for PowerHA® SystemMirror® Enterprise Edition utilizes several different network types that you can configure in PowerHA® SystemMirror®.

You can configure the following network types using GLVM for PowerHA® SystemMirror® Enterprise Edition in PowerHA® SystemMirror®.

XD_data

An IP-based network used by geographically mirrored volume groups in a PowerHA® SystemMirror® cluster for transferring the data between the RPV devices. This network is also used for participation in RSCT protocols and heartbeating. You can have up to four such networks. To increase high availability and performance, have two or more data mirroring network configured in the cluster. (All such networks should have similar bandwidth and latency characteristics, to enable load balancing for RPV client-server communication).

If one of the data mirroring networks fails, the GLVM data mirroring can continue over the redundant networks. Also, you have the flexibility to configure several low bandwidth XD_data networks and take advantage of the aggregate network bandwidth (of course, you can combine high bandwidth networks in the same manner).

You may need to check with your network provider the characteristics of the provisioned IP network. If the provider does not implement this network with multiple physical networks underneath, request to have an additional IP network to configure it as an additional XD_data network.

XD_ip

An IP-based network used for participation in RSCT protocols, heartbeating, and client communication. Typically, the network used for the XD_ip network type does not have enough bandwidth or is too slow to be suitable for data mirroring.

You should have more than one data mirroring network configured, and additional heartbeating networks in a GLVM for PowerHA® SystemMirror® Enterprise Edition cluster. Heartbeating enables PowerHA® SystemMirror® to distinguish between a genuine remote site outage and a failure of an XD_data network.

Planning for IP address takeover on XD networks

IP Address Takeover via IP aliases is the default method for IP label and address recovery in PowerHA® SystemMirror®, and it is enabled in GLVM for PowerHA® SystemMirror® Enterprise Edition by default.

XD_data networks with IP Address Takeover via IP aliases must have a service IP label that falls over between nodes in the cluster. Clients outside of the cluster can use this service IP label. This IP address is associated with the primary instance of a resource group, and can be moved via aliasing between nodes at the same site or across sites.

Since an IP address at one site may not be valid at the other site because of subnet issues, you can configure site-specific service IP labels to handle this situation.

Site-specific service IP labels are only activated when their resource group is online primary on the specified site. The service IP label can fall over to other nodes on the specified site. Except for the site-specific limitation, these service IP labels have all the same functions as regular service IP labels, except the ability to work with NFS cross-mounts between the sites.

For more information on configuring networks and IP labels, see *Administering PowerHA® SystemMirror®*.

Planning for XD_data Networks

An XD_data network type is used by geographically mirrored volume groups in a PowerHA® SystemMirror® cluster for transferring the data between the RPV devices. To plan for the XD_data networks, consider the following.

- In GLVM for PowerHA® SystemMirror® Enterprise Edition, you can have up to four XD_data networks in the cluster.
- Estimate the required network bandwidth for each XD_data network.

- IP Address Takeover via IP aliases is enabled by default for XD_data networks. Plan to assign persistent node IP labels on each node that is part of the network.
Include an alias service IP label for clients to access and decide whether to make this a site-specific service IP label.

Note: If you do not want any IP Address Takeover, you must convert the network type to private.

Planning for the XD_ip Networks

An XD_ip network type is used for heartbeating between nodes and across sites. The XD_ip network has IP Address Takeover via IP aliases enabled by default. You must assign an alias IP service label on the network. You can decide whether to make this a site-specific service IP label.

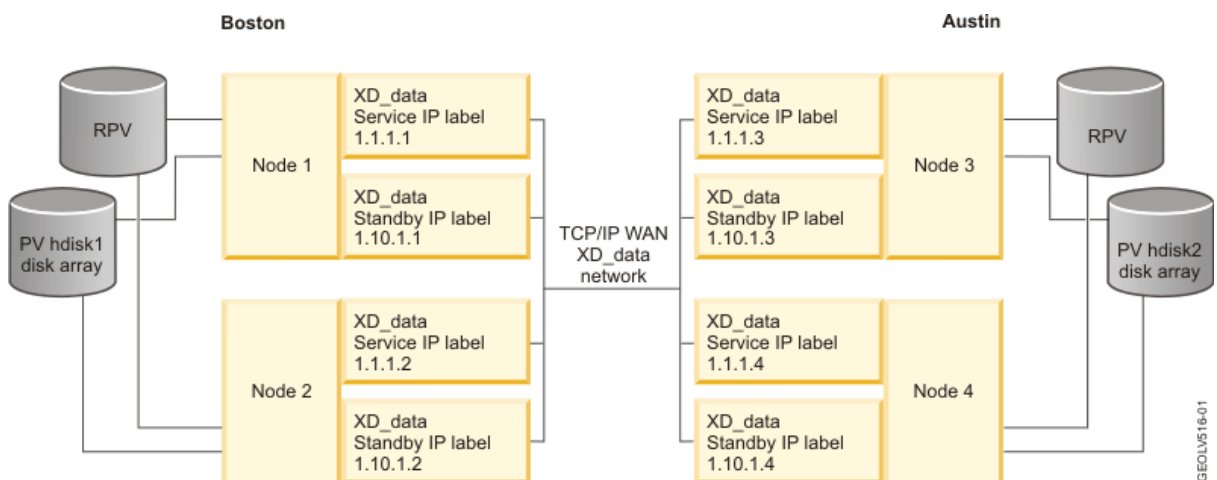
Establishing the connection over the XD_data network

When an RPV client is in an available state on a node, it uses the persistent IP or service IP label of the XD_data network on that node as its outgoing IP address for connecting to the RPV server.

- GLVM for PowerHA® SystemMirror® Enterprise Edition determines the IP address of the RPV server address based on the knowledge of which node at the remote site is currently hosting the mirror copy of the resource group. If the mirror copy of the resource group moves to another node at the remote site, GLVM for PowerHA® SystemMirror® Enterprise Edition dynamically reconfigures the IP address of the RPV server during the event processing.
- The RPV server also needs to know which IP addresses from the associated RPV client are allowed to initiate connections issued from the nodes that host the primary copy of the resource group that is geographically mirrored.
- For this purpose, when you configure the RPV server, it lets you specify a list of possible RPV client IP addresses it can accept. This way, if during the cluster events (such as a failover) the mirror copy of the resource group moves to another node on the remote site and the client IP address changes, the RPV servers recognize this address in the list of allowed IP addresses of the RPV client for this XD_data network.

Sample cluster configuration with one XD_data network using node-bound IP labels

The following figure shows how you can configure the RPV client and server IP addresses in PowerHA® SystemMirror® in a cluster with two nodes at each site. The diagram displays only one XD_data network, though it is recommended to have more:



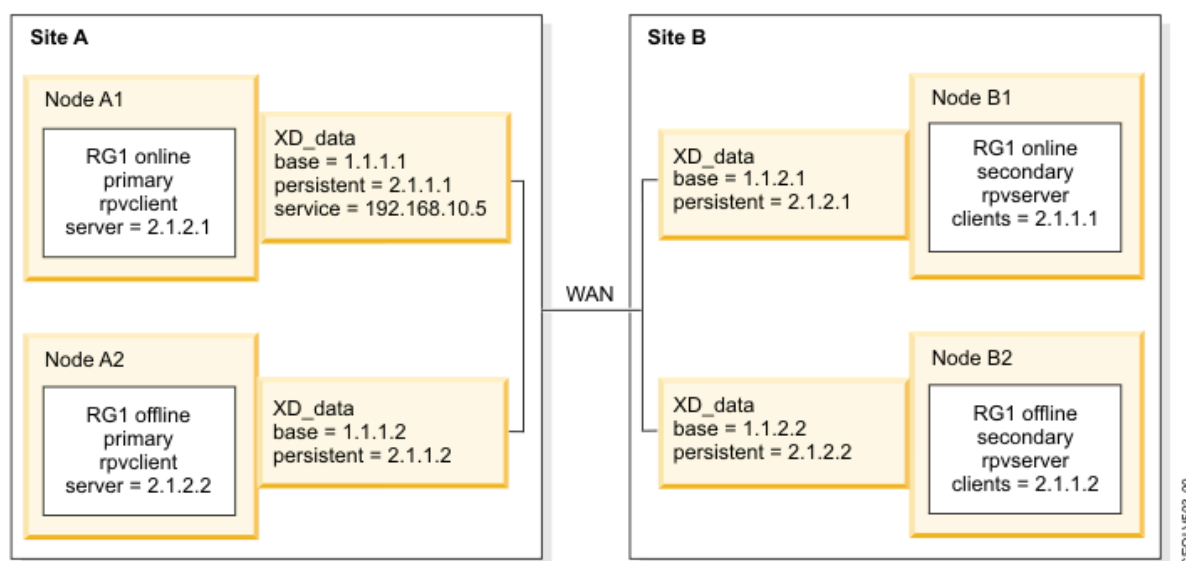
Understanding the connection over two XD_data networks

If a cluster has two data mirroring networks, the data mirroring communication occurs over both networks.

- When RPV clients from Node 1 (Boston) communicate with RPV servers on Node 3 (Austin), they use both addresses from both networks on Node 1 to connect to both addresses at both networks on Node 2. Note, in the two networks configuration, you must define twice as many IP addresses as with a configuration that has one XD_data network.
- If one of the XD_data networks fails, the RPV clients from Node 1 (Boston) communicate with RPV servers on Node 3 (Austin) by using just the addresses from the remaining functioning XD_data network. After you restore the failed network, GLVM for PowerHA® SystemMirror® Enterprise Edition automatically resumes using both networks for RPV client server communication.

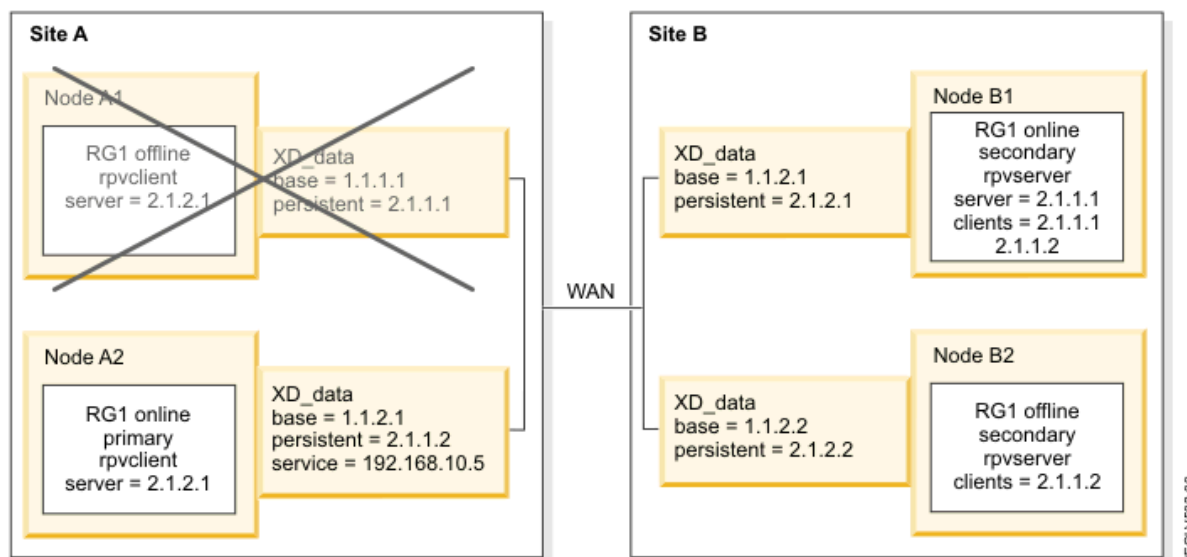
Sample configurations with IP address takeover via IP aliases on the XD_data network

The following figure is an example of what will happen during a fallover of a GLVM client when IP Address Takeover via IP aliases is used. The base, persistent, and service IP addresses, and the IP address the RPVclients and RPVserver communicate with are shown. For simplicity, the RPVclient and RPVserver configuration to handle site failure is not shown. Although it is strongly recommended to use the XD_ip network for additional service IP labels, the XD_data network is used in the example.

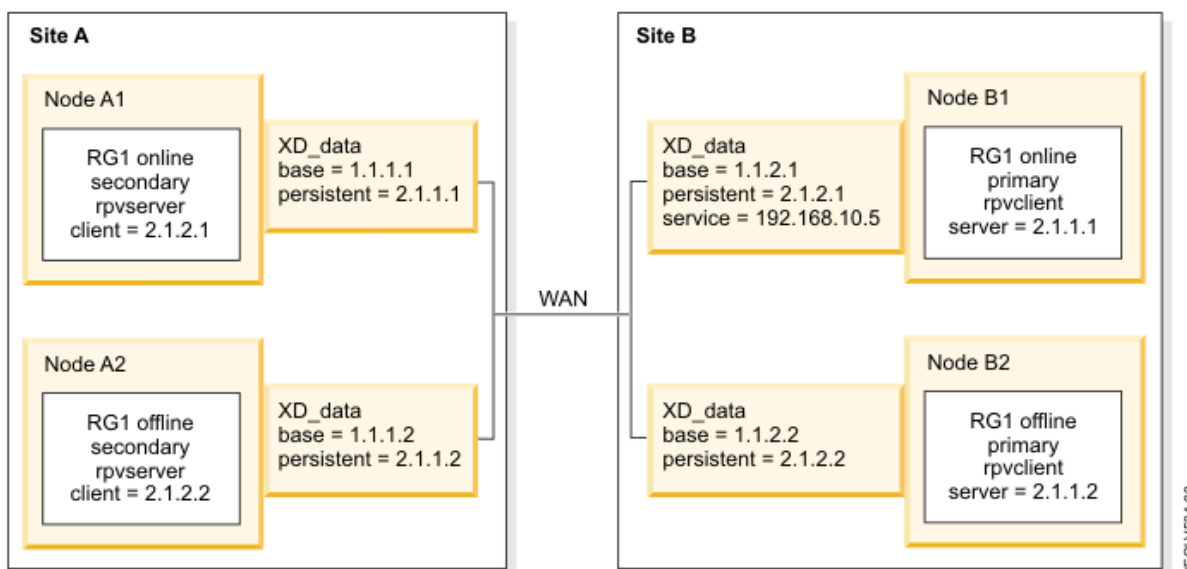


In the initial configuration of the GLVM cluster, the RPVclient for node A1 is bound to the persistent IP address 2.1.1.1 and uses the RPVserver on node B1, which is bound to persistent IP address 2.1.2.1. The corresponding setup is true for nodes A2 and B2 where the RPVclient on node A2 uses the RPVserver on node B2.

The following figure shows the GLVM configuration after resource group RG1 falls over from node A1 to node A2. The RPVclient on node A2 has been updated to use the RPVserver on node B1, which has had its client list updated to include the persistent IP address 2.1.1.2 on node A2. Also, the service IP label 192.168.10.5 now resides on the XD_data interface on node A2.



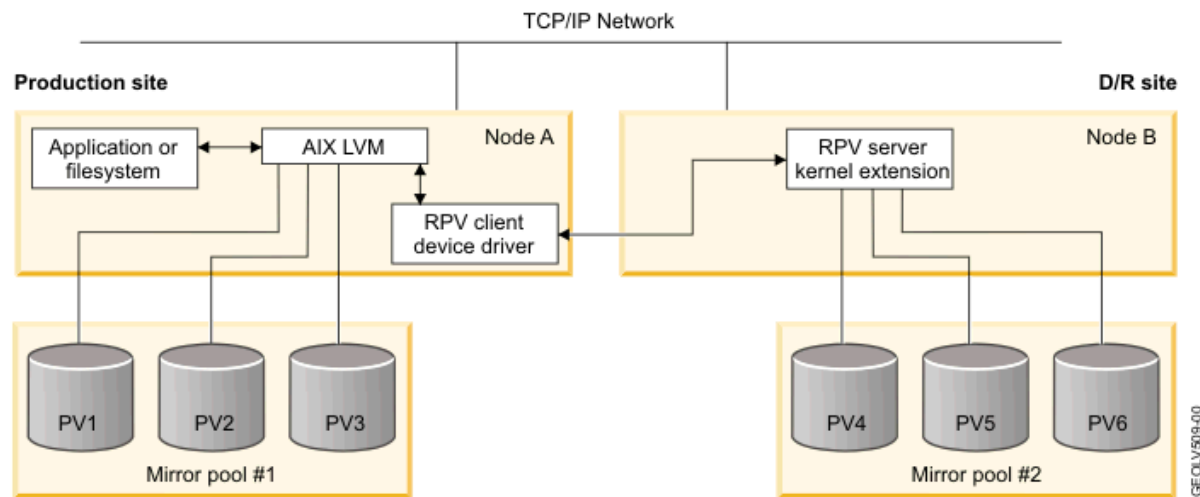
The following figure shows the GLVM configuration after the resource group falls over to the other site and remains active there. At this point, the failed nodes at the failed site have rejoined the cluster.



Planning for asynchronous mirroring

The asynchronous mirroring function is split into several functional areas.

The following figure displays asynchronous mirroring at the high level, based on a simple two-node stand-alone GLVM configuration. In this example, Node A resides at the production site and Node B resides at the disaster recovery site, as shown in the following diagram:



Mirror pools

Mirror pools provide a way for you to group disks together within a volume group.

GLVM requires that volume groups use super strict mirror pools. Super strict mirror pools must follow these rules:

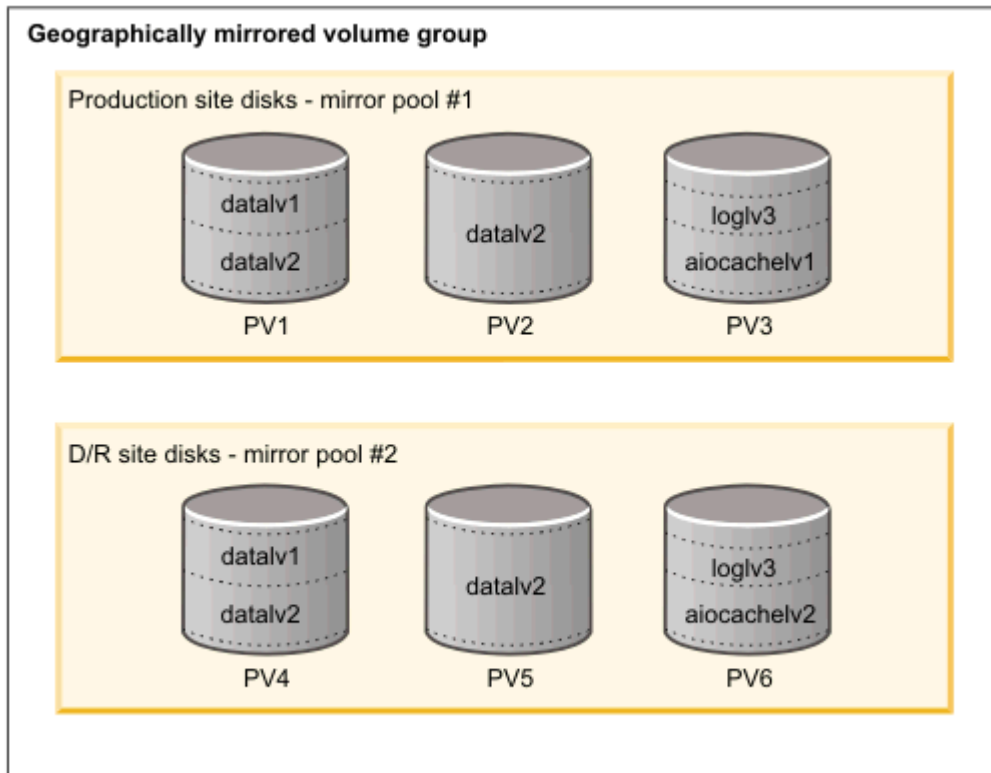
- Local and remote disks cannot belong to the same mirror pool.
- No more than three mirror pools per volume group.
- Each mirror pool must contain at least one copy of each logical volume, with some exceptions:
 - When you create a logical volume, you need to configure it so that each mirror pool gets a copy. However, if you create a mirror pool in a volume group where logical volumes already exist, logical volume copies are not automatically created in the new mirror pool. You need to create them by running the **mirrorvg** command or the **mklvcopy** commands.
 - Asynchronous GLVM mirroring will require a new type of logical volume for caching of asynchronous write requests. This logical volume should not be mirrored across sites. Super strict mirror pools handle this new **aio_cache** logical volume type as a special case.

Additionally, mirror pools provide extra benefits to the asynchronous mirroring function:

- Synchronous or Asynchronous is an attribute of a mirror pool. Rather than having to configure individual RPV devices, mirror pools provide a convenient way for users to manage asynchronous mirroring at a higher level.
- The decision of whether to mirror synchronously or asynchronously is made at the mirror pool level. Therefore, you can decide to mirror from the production site to the disaster recovery site asynchronously, and then mirror from the disaster recovery site back to the production site synchronously. This can be accomplished by configuring the mirror pool that contains the disaster recovery site disks as asynchronous while configuring the mirror pool that contains the production site disks as synchronous.

This example shows a geographically mirrored volume group, where the disks at the production site have been placed into Mirror Pool 1, and the disks at the disaster recovery site have been placed into Mirror Pool 2. Both mirror pools are configured for asynchronous mirroring.

The following figure displays a geographically mirrored volume group:



GEOLV512-00

This volume group has a total of five logical volumes. The user data is stored in three logical volumes. Logical volumes *datalv1* and *datalv2* contain file systems, and logical volume *loglv3* contains the file system log. These three logical volumes are mirrored across both sites, because they have copies in both mirror pools.

Logical volumes *aiocachelv1* and *aiocachelv2* are used for caching of asynchronous write requests. They are not mirrored across both sites.

In the figure, the volume group is varied online at the production site. Writes to the local disks in Mirror Pool 1 occur synchronously, in the usual manner. However, writes to the remote disks in Mirror Pool 2 are processed asynchronously. The RPV clients on Node A use logical volume *aiocachelv1*, which resides on a local disk in Mirror Pool 1, for caching of asynchronous write requests.

The opposite would be true if the volume group was varied online at the disaster recovery site. Writes to the local disks in Mirror Pool 2 would occur synchronously, in the usual manner. However, writes to the remote disks in Mirror Pool 1 would be processed asynchronously. The RPV clients on Node B would use logical volume *aiocachelv2*, which resides on a local disk in Mirror Pool 2, for caching of asynchronous write requests.

Asynchronous mirroring best practices

In order to make asynchronous GLVM configurations as highly available and efficient as possible, there are several choices you need to understand.

Protecting against disk failure

AIX® LVM data mirroring is designed to provide protection against disk and disk adapter failures. However, some asynchronous GLVM configurations may not always provide complete protection against these failures. You can avoid this problem by adding more disk redundancy.

For example, if there is one mirror copy of the data at each site, the mirror copy at the disaster recovery site is likely to be back-level most of the time, because the remote mirroring constantly lags behind. If a hardware failure results in a loss of disks containing portions of the production site mirror copy and the cache of remote physical volume writes, then the possibility exists that the only surviving copy of some logical partitions may be the back-level physical partitions that belong to the disaster recovery site mirror copy. Unfortunately, asynchronous mirroring lacks the full redundancy that you normally get from having two LVM mirror copies of the data.

One solution is to configure asynchronous GLVM mirroring with two mirror copies of the data at the production site. This makes the production site data more highly available and less likely that the last surviving copy of a logical partition will be stale. However, this is not a perfect solution. LVM allows only three mirror copies. If you create two mirror copies of the data at the production site, you are allowed to have only one mirror copy at the disaster recovery site. This configuration is still vulnerable to disk failures at the disaster recovery site, which may leave you without a complete copy of the data in the event of a production site disaster. Additionally, if you choose **Online On Either Site** as the PowerHA® SystemMirror® site policy, then data mirroring from the disaster recovery site back to the production site requires updating two remote copies, requiring twice as much network bandwidth. You can achieve better performance using this method by selecting **Prefer Primary Site** as the PowerHA® SystemMirror® site policy. Using this policy, data mirroring always goes from the production site to the disaster recovery site, where there is only one mirror copy. If the application workload runs at the disaster recovery site for a while, the data on the production site disks becomes stale. When the production site is recovered and brought back into the cluster, the **Prefer Primary Site** policy causes the application workload to move back to the production site. From there, the **syncvg** command can more efficiently bring the two production site mirror copies up to date because it only needs to read from one remote copy in order to update two local copies. If the application workload had remained at the disaster recovery site, then the **syncvg** command would have to write to two remote copies.

An even better solution is to combine asynchronous GLVM mirroring with a disk subsystem that has built-in data mirroring or RAID capabilities at both sites. This makes the data at both sites more highly available. This provides better protection at both sites and it reduces the chances that LVM or the RPV device driver will experience I/O errors.

Protecting against site failure

You need to protect your site from accessing back-level data if the production site node should crash.

If there is only one node at the production site and you configure PowerHA® SystemMirror® to automatically allow access to back-level data upon failover to the disaster recovery site, then your configuration is vulnerable to data divergence should the production site node crash. This vulnerability can be reduced, but not completely eliminated, by adding a second node to the production site. In this case, if the original production site node crashes, PowerHA® SystemMirror® can perform a local peer failover to move the application workload over to the other production site node. This would avoid data divergence. However, the configuration would still be vulnerable to an outage that causes both production site nodes to suddenly fail.

Avoiding risky data divergence recovery

After data divergence has occurred, running the application workload at the disaster recovery site for a long time, and then deciding to back out those updates and go back to the production site version of the data may pose a risk for data loss.

The ability to keep the production site updates and throw away the disaster recovery site updates is intended for cases where the user intervenes shortly after a site failover occurs, before there is a chance to make many updates to the disks at the disaster recovery site. When choosing to go back to the production site version of the data, the user may want to protect against data loss by running the **savevg** command to back up the disaster recovery site's copy of the volume group, so that it can get it back if something goes wrong.

Installing GLVM for PowerHA® SystemMirror® Enterprise Edition

These topics describe how to install GLVM for PowerHA® SystemMirror® Enterprise Edition.

Installation components

The GLVM for PowerHA® SystemMirror® Enterprise Edition software is comprised of the several different components.

Some of the components that are required for the installation of GLVM are now available as part of the AIX® operating system. The changes to the packaging of the installation components are for PowerHA® SystemMirror® Enterprise Edition 7.2.2, or later. Review the following table and verify that the filesets are installed before you configure GLVM.

The GLVM for PowerHA® SystemMirror® Enterprise Edition software is comprised of the following components:

| Component Name | Description | Content | Filesets |
|----------------------------------------------------|----------------------------------------------------------------------------------------------|---------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GLVM | Builds upon AIX® to provide the geographic mirroring and storage management functions. | Includes filesets for the GLVM utilities and the RPV device driver. | <ul style="list-style-type: none"> glvm.rpv.util glvm.rpv.client glvm.rpv.server <div style="border: 1px solid blue; padding: 5px; margin-top: 10px;"> <p>Note: These filesets are on the AIX® 7.2.2, or later, or AIX® 7.1.5, or later, base media.</p> </div> |
| GLVM for PowerHA® SystemMirror® Enterprise Edition | Provides integration of the PowerHA® SystemMirror® high availability function with the GLVM. | Includes filesets for GLVM for PowerHA® SystemMirror® Enterprise Edition. | <ul style="list-style-type: none"> cluster.xd.license cluster.xd.glvm |

Installation prerequisites

Before you install GLVM for PowerHA® SystemMirror® Enterprise Edition, make sure that the necessary software is installed on the cluster nodes.

You must have the following software installed before you can install GLVM for PowerHA® SystemMirror® Enterprise Edition.

- Latest version of PowerHA® SystemMirror®, AIX®, and Reliable Scalable Cluster Technology (RSCT). RSCT is typically installed as part of the AIX® operating system.
- The PowerHA® SystemMirror® Enterprise Edition software uses about 1 MB of disk space. Ensure that the /usr file system has at least 1 MB free disk space for the upgrade.
- You must install the following filesets that are on the AIX® 7.2.2, or later, or AIX® 7.1.5, or later, base media:
 - glvm.rpv.util
 - glvm.rpv.client
 - glvm.rpv.server

Perform the installation process as the root user.

Installation options

You can install GLVM in various ways, depending on what software and hardware configuration you have before installing GLVM for PowerHA® SystemMirror® Enterprise Edition:

The following list includes possible installation options for installing GLVM for PowerHA® SystemMirror® Enterprise Edition.

- Installing GLVM for PowerHA® SystemMirror® Enterprise Edition on the working cluster that is running PowerHA® SystemMirror®:
Before the installation, you may want to ensure that you have two (no more) sites defined to the cluster, all nodes at each location belong to the corresponding site and the inter-site management policy is set to nonconcurrent. Optionally, you can make these changes after you install GLVM for PowerHA® SystemMirror® Enterprise Edition.
- Installing GLVM for PowerHA® SystemMirror® Enterprise Edition together with installing PowerHA® SystemMirror®:
This type of installation entails planning for the initial PowerHA® SystemMirror® cluster setup with two sites, along with planning for the PowerHA® SystemMirror® Enterprise Edition for GLVM disaster recovery solution between two sites.

Filesets on the installation media

There are two filesets on the installation media that are required to install GLVM for PowerHA® SystemMirror® Enterprise Edition.

The following filesets are required to install GLVM for PowerHA® SystemMirror® Enterprise Edition:

cluster.xd.glvm

GLVM for PowerHA® SystemMirror® Enterprise Edition RPV Support

cluster.xd.license

PowerHA® SystemMirror® Enterprise Edition Feature License Agreement

Installation methods

Install the GLVM for PowerHA® SystemMirror® Enterprise Edition software on each cluster node (server).

You can install GLVM for PowerHA® SystemMirror® Enterprise Edition software using one of the following methods:

- From an installation server
- From a hard disk to which the software has been copied
- From the installation media

Installing from an installation server

You can create a GLVM for PowerHA® SystemMirror® Enterprise Edition installation server (containing the GLVM for PowerHA® SystemMirror® Enterprise Edition software images for installation) on one node and then load the images onto the remaining cluster nodes. This is the fastest way to install GLVM for PowerHA® SystemMirror® Enterprise Edition.

GLVM for PowerHA® SystemMirror® Enterprise Edition supports the Network Installation Management program and Alternate Disk Migration.

Installing from a hard disk

To install GLVM for PowerHA® SystemMirror® Enterprise Edition software from your hard disk, copy the software from the installation media to the hard disk prior to installation

To copy GLVM for PowerHA® SystemMirror® Enterprise Edition software from your hard disk :

1. Insert the PowerHA® SystemMirror® Enterprise Edition CD into the CD-ROM drive.
2. Enter `smit bffcreate`
The **Copy Software to Hard Disk for Future Installation** panel appears.
3. Enter the name of the CD-ROM drive in the **INPUT device / directory for software** field and press Enter. If you are unsure of the input device name, press F4 to list available devices. Select the correct drive and press Enter. That drive name appears in 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 field values as follows:

| <i>Copy Software to Hard Disk for Future Installation</i> | |
|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| SOFTWARE name | Press F4 for a software listing. Install the images for GLVM for PowerHA® SystemMirror® Enterprise Edition. For a list of the GLVM for PowerHA® SystemMirror® Enterprise Edition images, see the section Contents of the installation media. |
| DIRECTORY for storing software | Change the value to the storage directory that all nodes using GLVM for PowerHA® SystemMirror® Enterprise Edition can access. |

6. Enter values for the other fields as appropriate for your site.

7. When you are satisfied with the entries, press Enter.
SMIT responds: **Are you sure?**
8. Press Enter again to copy the software.

Result

Installing GLVM for PowerHA® SystemMirror® Enterprise Edition from the Hard Disk

After the GLVM for PowerHA® SystemMirror® Enterprise Edition software is on your system, follow the instructions in the section Installing the GLVM for PowerHA® SystemMirror® Enterprise Edition software from the installation media to install the software.

Installing from the installation media

If you install the GLVM for PowerHA® SystemMirror® Enterprise Edition software from the installation media, you must install the software directly onto each cluster node.

To install the GLVM for PowerHA® SystemMirror® Enterprise Edition software on a server node, complete the following steps.

1. Insert the installation media into the device the reads the installation media.
2. Enter `smit install_all`
SMIT displays the first **Install and Update from ALL Available Software** panel.
3. Enter the device name of the installation media 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 install directory, press F4 to list the available devices. Then select the correct device or directory and press Enter. The correct value is entered into the **INPUT device / directory** field as the valid input device.
4. Enter field values as follows. Press F1 for help on any field.

Use F4 to list the software before proceeding with the installation. This way you can install the message catalogs.

| <i>INPUT device / directory for software fields</i> | |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| INPUT device / directory for software | This field shows the device or directory you specified earlier. |
| SOFTWARE to install | Press F4 for a software listing. In the software list, use the arrow keys to locate all software filesets associated with an image. For a list of filesets, see the section Contents of the installation media. Next press F7 to select either an image or a fileset. Then press Enter after making all selections. Your selections appear in this field. |
| PREVIEW only? | If set to yes , the preview option checks and verifies that installation prerequisites are met, for instance that required software is installed and sufficient disk space is available. Press F1 for details. When you are ready to perform the actual installation, set this field to no . |
| COMMIT software updates? | This field applies only when you install software updates (PTFs). See the F1 help for details. |
| SAVE replaced files? | This field applies only when you install software updates (PTFs). If you select no in response to commit software updates? select yes for this field. See F1 help for details. |
| AUTOMATICALLY install requisite software | Set this field to no if the prerequisite software 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. See F1 help for details. |

| Field | Value |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EXTEND filesystems if space needed? | Select yes if you have adequate hard disk space, no if you have limited space. See F1 help for details. |
| OVERWRITE same or newer versions? | For new installations, leave this field set to no . Set it to yes if you are reinstalling the software. If you set this field to yes , set the Automatically install requisite software field to no . See the F1 help for details. |
| VERIFY install and check file sizes? | Select yes if you want the system to perform some checks on the software you installed. See F1 help for details. |
| 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. See F1 for information. |
| ACCEPT new license agreements? | Select yes for this item to proceed with installation. If you choose no , installation may stop with a warning that one or more filesets 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. |

- When you are satisfied with the entries, press Enter. SMIT responds **Are you sure?**
- Press Enter to install the software.

Result

Review the GLVM for PowerHA® SystemMirror® Enterprise Edition release notes in the `/usr/es/sbin/cluster/release_notes_xd` file for information that does not appear in the product documentation.

Verifying the installed cluster definition

After the GLVM for PowerHA® SystemMirror® Enterprise Edition software is installed on all nodes, you should verify the configuration. Verification provides errors or warnings to ensure that the cluster definition is the same on all nodes.

Note: You can verify the installed cluster definition if the software is being installed on an existing cluster. You cannot verify the installed cluster definition on a new cluster because the cluster topology has not been defined yet.

To verify the cluster, complete the following steps.

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

Result

Verification checks both the PowerHA® SystemMirror® and the GLVM for PowerHA® SystemMirror® Enterprise Edition configuration.

Entries in the `/etc/services` file

After GLVM for PowerHA® SystemMirror® Enterprise Edition is installed, an entry is made in the `/etc/services` file.

The following entry for the port number and connection protocol is automatically added to the `/etc/services` file on each node on which you installed the software:

```
rpv 6192/tcp
```

This defines the TCP/IP port number which is to be used for RPV client-server communication.

Addressing problems during the installation

If you experience problems during the installation, the installation program usually performs a cleanup process automatically.

If the cleanup process is not automatically started, complete the following steps.

1. Enter `smit install` to display the Installation and Maintenance menu.
2. Select **Software Maintenance and Utilities**.
3. Select **Clean Up After a Failed or Interrupted Installation**.
4. Review the SMIT output (or examine the `/smit.log` file) for the cause of the interruption.
5. Fix any problems and repeat the installation process.

Upgrading GLVM for PowerHA® SystemMirror® Enterprise Edition

You can upgrade GLVM from prior versions of PowerHA® SystemMirror® and PowerHA® SystemMirror® Enterprise Edition. The supported migration paths depend on the currently installed version and the version you are migrating too.

Upgrade prerequisites

Before you upgrade:

- Verify that all PowerHA® SystemMirror® cluster nodes are running the latest version of PowerHA® SystemMirror®.
- Verify that you are familiar with the process for installing and configuring PowerHA® SystemMirror® Enterprise Edition. Install GLVM from the [Installation media](#).
- Verify that your system meets the installation prerequisites, see [Installation prerequisites](#).
- Verify that each cluster node has its own PowerHA® SystemMirror® Enterprise Edition license.
- Perform the installation process as the root user.

Verifying the upgraded cluster definition

After the GLVM for PowerHA® SystemMirror® Enterprise Edition software is installed on all nodes, verify the configuration. Verification provides errors or warnings to ensure that the cluster definition is the same on all nodes.

To verify the cluster:

1. Enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Verify and Synchronize Cluster Configuration (Advanced)**, and press Enter.
Verification verifies both PowerHA® SystemMirror® and GLVM for PowerHA® SystemMirror® Enterprise Edition configuration.

Result You cannot synchronize a mixed-version cluster. When upgrading, you should not leave the cluster at mixed versions of PowerHA® SystemMirror® for long periods of time. New functions supplied with the latest PowerHA® SystemMirror® version are available only when all nodes have been upgraded and the cluster has been synchronized. Do not expect commands such as `clfindres` to supply correct information in a mixed cluster.

Recovering from a failed installation

If your installation fails, there are several different steps you need to take to recover.

When you install PowerHA® SystemMirror® Enterprise Edition, the `cl_convert` command runs automatically to convert the PowerHA® SystemMirror® configuration database from a previous release to that of the current release. If the installation fails, run `cl_convert` from the command line to convert the database.

In a failed conversion, run `cl_convert` using the `-F` flag.

To run a conversion utility requires:

- The root user privilege
- The PowerHA® SystemMirror® Enterprise Edition version from which you are converting.

The **cl_convert** utility logs conversion progress to the **/tmp/clconvert.log** file so that you can gauge conversion success. This log file is regenerated each time **cl_convert** or **clconvert_snapshot** runs.

For more information about **cl_convert** and **clconvert_snapshot**, see the respective man pages or see Administering PowerHA® SystemMirror®.

Modifying previous cluster snapshots

After you upgrade your PowerHA® SystemMirror® Enterprise Edition software, you might want to restore one or more of the previous version cluster snapshots you created using the Cluster Snapshot utility.

The default directory path for storage and retrieval of a snapshot is **/usr/es/sbin/cluster/snapshots**; however, you may have specified an alternate path using the **SNAPSHOTPATH** environment variable. Check in these locations before using the **/usr/es/sbin/cluster/conversion/clconvert_snapshot** utility to convert the snapshot.

The snapshot is based on your full PowerHA® SystemMirror® configuration, including the configuration with replicated resources configured with GLVM for PowerHA® SystemMirror® Enterprise Edition. The **clconvert_snapshot** utility updates PowerHA® SystemMirror® configuration data with new information for the latest version.

To convert and apply a cluster snapshot, enter:

```
clconvert_snapshot -v version# -s snapshot_file_name
```

where the **-s** flag is used with the snapshot filename you want to update or apply, and the **-v** flag is used with the version of the saved snapshot.

Configuring geographically mirrored volume groups

These topics describe how to configure geographically mirrored volume groups, their corresponding logical volumes, and remote physical volumes (RPVs). Having these entities set up allows you to have a copy of data for your applications mirrored at a remote site with the GLVM for PowerHA® SystemMirror® Enterprise Edition support.

Configuration prerequisites

Before configuring a GLVM environment, you will need to ensure some prerequisite steps have been taken.

Ensure that you do the following:

- Install AIX® on all nodes in the cluster.
- Configure clusters that include site definitions.
- Service IP addresses on the remote site addresses serve as service IP labels and IP addresses for the RPV servers (that reside on the remote site). You will make the addresses known to the local site.
- Using the standard AIX® LVM SMIT panels, configure volume groups, logical volumes and file systems for which you plan to configure geographic mirrors with the GLVM utilities. Ensure that either standard or enhanced concurrent mode LVM volume groups exist for the data you plan to be geographically mirrored.

Note: If you want to turn your existing enhanced concurrent volume groups into geographically mirrored volume groups by adding RPVs, you must have PowerHA® SystemMirror® cluster services running on the nodes. (Although you can create standard volume groups in AIX® and add RPVs to these groups using the GLVM utilities in SMIT, to add RPVs to enhanced concurrent volume groups, PowerHA® SystemMirror® cluster services must be running).

- For all logical volumes that are planned to be geographically mirrored, ensure that the inter-disk allocation policy is set to Super Strict.

- Create volume groups and mirror pools for use with asynchronous mirroring. Mirror pools are required for using asynchronous mirroring but are optional when using synchronous mirroring.

GLVM Configuration Assistant

You can use the GLVM Configuration Assistant to automatically create and configure eligible volume groups for Geographical Logical Volume Manager (GLVM) mirroring.

GLVM Configuration Assistant supports both Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) addresses for persistent IP labels.

The GLVM Cluster Configuration Assistant creates informative messages to indicate what processes are being performed. The messages are logged in the `clutils.log` file that is in the `/var/hacmp/log` directory.

You can use the GLVM Cluster Configuration Assistant to create a geographically mirrored volume group that has the following characteristics:

- The geographically mirrored volume group exists on all nodes in the cluster.
- All available persistent labels on XD_data networks are used.
- A resource group is created with the following attributes:

Prefer Primary Site

The primary instance of the resource group is brought ONLINE during startup on the primary site. The secondary instance is started on the other site. The primary instance falls back when the primary site rejoins the cluster.

Online On Home Node Only

The resource group is activated on the first participating node that becomes available. The settling time for resource groups is used only for the startup policy of this resource group.

Fallover To Next Priority Node In The List

The resource group that is online on only one node at a time follows the default node priority order that is specified in the resource group's node list.

Never Fallback

The resource group does not fall back when a higher priority node joins the cluster.

Prerequisites for using the GLVM Configuration Assistant

To use the GLVM Configuration Assistant, you must install the correct filesets and you must configure your environment correctly.

Consider the following prerequisites for using the GLVM Configuration Assistant:

- A cluster is configured with sites.
- A repository disk is defined in the cluster configuration.
- The verification and synchronization process completes successfully on the cluster.
- XD_data networks with persistent IP labels are defined on the cluster.
- The network communication between the local site and remote site is working.
- All PowerHA® SystemMirror® services are active on both nodes in the cluster.
- The `/etc/hosts` file on both sites contains all of the host IP, service IP, and persistent IP labels that you want to use in the GLVM configuration.
- PowerHA® SystemMirror® 7.2.2, or later, is installed on all nodes in the cluster.
- Verify that the remote site has enough free disks and enough free space on those disks to support all of the local site volume groups that are created for geographical mirroring.
- AIX® 7.2.2, or later, or AIX® 7.1.5, or later, is installed on all nodes in the cluster. Although previous versions of the AIX® operating system are supported, it is recommend that you use PowerHA® SystemMirror® 7.2.2 with AIX® 7.2.2, or later, or AIX® 7.1.5, or later.

- The following filesets must be installed on your system:
 - `cluster.xd.license`
 - `cluster.xd.glv`
 - `glvm.rpv.util`
 - `glvm.rpv.client`
 - `glvm.rpv.server`

Other:

- The `cluster.xd.license` and `cluster.xd.glv` filesets are on the PowerHA® SystemMirror® media. The `glvm.rpv.util`, `glvm.rpv.client`, and `glvm.rpv.server` filesets are on the AIX® 7.2.2, or later, or AIX® 7.1.5, or later, base media.
- To use the GLVM Configuration Assistant, you must install Python version 2.7.10, or later.

Configuring an asynchronous geographically mirrored volume group by using the GLVM Configuration Assistant

Asynchronous mirroring updates the local site immediately and the remote site is updated when bandwidth between sites is available.

To configure an asynchronous geographically mirrored volume group (GMVG) by using the GLVM Configuration Assistant, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Make Applications Highly Available (Use Smart Assist) > GLVM Configuration Assistant > Configure Asynchronous GMVG**, and press Enter.
3. Enter the field values as follows:

| Field | Description |
|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enter the name of the VG | Enter the name of the volume group that you want to create as a geographically mirrored volume group. If the resource group is created by using the GLVM Configuration Assistant, the volume group name is appended with <code>_RG</code> . For example, if the volume group name is <code>datavg1</code> , the resource group name is <code>datavg1_RG</code> |
| Select disks to be mirrored from the local site | Press F4 to display a list of available disks. Press F7 to select the disks that you want to geographically mirror from the local site. After all disks are selected, press Enter. |
| Select disks to be mirrored from the remote site | Press F4 to display a list of available disks. Press F7 to select the disks that you want to geographically mirror from the remote site. After all disks are selected, press Enter. |
| Enter the size of the ASYNCH cache | Enter the number of physical partitions (PPs) on the volume group. The number that you enter depends on the load of the applications and bandwidth available on the network. You might |

| Field | Description |
|-------|----------------------------------------------------------------|
| | need to enter different values for peak workload optimization. |

- Verify that all fields are correct and press Enter to create the asynchronous geographically mirrored volume group.
- Run the **lsvg** command to verify that the new volume group was created successfully.
- Run the **clshowres** command to verify that the new resource group was created successfully.

Configuring an synchronous geographically mirrored volume group by using the GLVM Configuration Assistant

Synchronous mirroring writes data to both the local and remote sites at the same time. Synchronous mirroring also keeps your disaster recovery site up to date.

To configure an synchronous geographically mirrored volume group (GMVG) by using the GLVM Configuration Assistant, complete the following steps:

- From the command line, enter **smit sysmirror**.
- From the SMIT interface, select **Cluster Applications and Resources > Make Applications Highly Available (Use Smart Assist) > GLVM Configuration Assistant > Configure Asynchronous GMVG**, and press Enter.
- Enter the field values as follows:

| Fields | Description |
|--------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enter the name of the VG | Enter the name of the volume group that you want to create as a geographically mirrored volume group. <div style="border: 1px solid blue; padding: 5px; margin-top: 10px;"> <p>If the resource group is created by using the GLVM Configuration Assistant, the volume group name is appended with <i>_RG</i>. For example, if the volume group name is <i>datavg1</i>, the resource group name is <i>datavg1_RG</i></p> </div> |
| Select disks to be mirrored from the local site | Press F4 to display a list of available disks. Press F7 to select the disks that you want to geographically mirror from the local site. After all disks are selected, press Enter. |
| Select disks to be mirrored from the remote site | Press F4 to display a list of available disks. Press F7 to select the disks that you want to geographically mirror from the remote site. After all disks are selected, press Enter. |

- Verify that all fields are correct and press Enter to create the synchronous geographically mirrored volume group.
- Run the **lsvg** command to verify that the new volume group was created successfully.
- Run the **clshowres** command to verify that the new resource group was created successfully.

GLVM DR Sizing tool

The performance of a volume group in an asynchronous Geographical Logical Volume Manager (GLVM) is affected by the bandwidth between sites and the cache configuration of a logical volume. For system requirement updates, you must monitor the network and performance of data replication of the GLVM volume group.

In PowerHA® SystemMirror® Version 7.2.7, the GLVM disaster recovery (DR) Sizing tool is used to monitor data replication and to provide suitable recommendations.

The GLVM DR Sizing tool performs the following functions:

- Collects the data writes and cache information for the local and remote sites.
- Analyzes the collected data.
- Provides recommendations to optimize the cache size and the network bandwidth.

Starting with PowerHA® SystemMirror® Version 7.2.8, and later, the GLVM DR Sizing tool can use the iPerf tool to provide suitable network latency recommendations. The iPerf tool has the following functions:

- Creates a server on the home node, with all the remote nodes connecting to this server.
- Transfer data packets on the network and collect network performance statistics.
- Provide information about network bandwidth and network latency.

Note: To get the network latency recommendations, you must use the iPerf tool version 2.0.9 on all virtual machines. For more information, see the [iPerf Tool](#).

The GLVM DR Sizing tool includes the **cl_survey** utility that is used to evaluate the current disk usage on the system. The **cl_survey** utility provides information to specific network that has sufficient bandwidth that can be used in an asynchronous GLVM volume group. The **cl_survey** utility works similar to UNIX® performance monitoring programs (for example, the **clperfstat** program). For more information, see the [cl_survey command](#).

The GLVM DR Sizing tool has the following limitations:

- An asynchronous GLVM must be configured in PowerHA® SystemMirror®, and the cluster services must be in the STABLE state.
- You cannot run the tool multiple times.
- Raw logical volumes are not supported.
- You must not change the configuration when the **cl_survey** tool is running. The **cl_survey** tool might stop if it detects any configuration changes when it is running.

GLVM Configuration Assistant enhancements

In PowerHA® SystemMirror® Version 7.2.5, or later, new enhancements are added to the Geographic Logical Volume Manager (GLVM) Configuration Wizard.

Monitoring remote physical volume statistics

The PowerHA® SystemMirror® GLVM Configuration Wizard collects the Remote Physical Volume (RPV) mirroring statistics and stores the information in the JavaScript™ Object Notation (JSON) format. You can use the RPV statistics by using any tool capable of displaying the JSON format. The RPV statistics data is automatically sent to the PowerHA® SystemMirror® GUI which displays it in a graphical format. For more information about capabilities of the PowerHA® SystemMirror® GUI, see [PowerHA® SystemMirror® graphical user interface \(GUI\)](#) topic.

When you use the GLVM Configuration Assistant, it creates an application monitor that collects the RPV mirroring statistics regularly. The application monitor is associated with an application controller, which is added to the PowerHA® SystemMirror® resource group that controls the GLVM replication.

After you start the cluster services, the application monitor automatically collects the RPV mirroring statistics and updates it in the JSON file. When the JSON file size increases more than 1 MB, the GLVM Configuration Assistant rotates the JSON file and stores up to eight copies of the numbered files with the oldest file being overwritten. The JSON file for each GLVM resource group is located in the `/var/hacmp/Log` subdirectory.

Note: In PowerHA® SystemMirror® Version 7.2.6, or later, the JSON file for each GLVM resource group is located in the `/var/hacmp/log/glvm` subdirectory.

You can view the JSON file or use it with other tools that display the JSON format. The JSON file name is constructed based on the resource group name. For example, if the resource group name is RG1, the JSON file name is `RG1_GLVM_Stat.JSON`. The `RG1_GLVM_Stat.JSON` file is located in the `/var/hacmp/log/glvm/RG1` directory. The `/var/hacmp/log/glvm/RG1` directory is added to the PowerHA® SystemMirror® File Collections to ensure that all JSON files are available on all cluster nodes during a resource group failover event.

Note: The application monitor triggers for every 5 minutes to collect the RPV mirroring statistics and to update it in the JSON file. The JSON file is updated only when the GLVM resource group is online.

Unconfiguring the GLVM components in the PowerHA® SystemMirror® cluster

You can remove the GLVM components from the PowerHA® SystemMirror® cluster along with associated resources by using the `theunconfigure` option.

GLVM resources that are removed include the synchronous and asynchronous volume groups and associated mirror pools, cache logical volumes, application controllers and monitors, and the RPV server and RPV clients. GLVM resource are removed from the local and remote sites through a single operation.

Note: To use the `theunconfigure` option for GLVM, the associated resource group must be in an offline state on all cluster nodes.

Prerequisites for using the GLVM Configuration Assistant

To use the GLVM Configuration Assistant, you must install the correct filesets and you must configure your environment correctly.

Consider the following prerequisites for using the GLVM Configuration Assistant:

- A cluster is configured with sites.
- A repository disk is defined in the cluster configuration.
- The verification and synchronization process completes successfully on the cluster.
- XD_data networks with persistent IP labels are defined on the cluster.
- The network communication between the local site and remote site is working.
- All PowerHA® SystemMirror® services are active on both nodes in the cluster.
- The `/etc/hosts` file on both sites contains all of the host IP, service IP, and persistent IP labels that you want to use in the GLVM configuration.
- PowerHA® SystemMirror® 7.2.2, or later, is installed on all nodes in the cluster.
- Verify that the remote site has enough free disks and enough free space on those disks to support all of the local site volume groups that are created for geographical mirroring.
- AIX® 7.2.2, or later, or AIX® 7.1.5, or later, is installed on all nodes in the cluster. Although previous versions of the AIX® operating system are supported, it is recommend that you use PowerHA® SystemMirror® 7.2.2 with AIX® 7.2.2, or later, or AIX® 7.1.5, or later.
- The following filesets must be installed on your system:
 - `cluster.xd.license`
 - `cluster.xd.glvm`
 - `glvm.rpv.util`

- `glvm.rpv.client`
- `glvm.rpv.server`

Other:

- The `cluster.xd.license` and `cluster.xd.glvm` filesets are on the PowerHA® SystemMirror® media. The `glvm.rpv.util`, `glvm.rpv.client`, and `glvm.rpv.server` filesets are on the AIX® 7.2.2, or later, or AIX® 7.1.5, or later, base media.
- To use the GLVM Configuration Assistant, you must install Python version 2.7.10, or later.

Configuring standard volume groups

If you choose to have standard volume groups, you can either first configure them in AIX® SMIT, add RPVs using GLVM utilities in SMIT, configure a PowerHA® SystemMirror® cluster and add geographically mirrored volume groups to resource groups or you can first configure a cluster without sites, add another site, and then add RPVs to existing volume groups, thus making them geographically mirrored volume groups.

To summarize, for standard volume groups, the configuration process is two-fold:

- Configuring RPV client-server pairs for a standard volume group at a node on each site
- Extending the geographically mirrored standard volume group to other nodes in the cluster

Configuring enhanced concurrent mode volume groups

If you choose to have enhanced concurrent volume groups that will also be geographically mirrored with GLVM, for some tasks you must have a PowerHA® SystemMirror® cluster configured and cluster services running.

To summarize, you can either create an enhanced concurrent volume group in AIX®, or take an existing enhanced concurrent volume group and extend it to another site. The high level overview of steps is included in these sections:

- Configuring RPV client-server pairs for an enhanced concurrent volume group on a node at each site
- Configuring RPVs, mirror copies and geographically mirrored logical volumes

Configuration overview for asynchronous mirroring

If you choose to use asynchronous mirroring, you need to use mirror pools, which in turn means that you need to use scalable volume groups.

Review the following information before you configure asynchronous mirroring:

- All disks in all mirror pools must be accessible to be configured for asynchronous mirroring.
- After a mirror pool is configured for asynchronous mirroring, some active disks are needed from each mirror pool to convert the mirror pool from asynchronous mirroring to synchronous mirroring. If you want to remove one or more mirror pools from a site that is down, disable asynchronous mirroring using the `chmp` command with the `-S` flag and `-f` flags.
- Asynchronous mirroring is only supported on nonconcurrent scalable volume groups with mirror pools set to be super strict.
- You must disable the auto-on and bad-block relocation options of the volume group.
- The volume group cannot be a snapshot volume group. The volume group cannot contain active paging-space logical volumes.
- The volume group must be varied on to make mirror pool changes.
- You cannot remove or reduce an `aio_cache` type logical volume once it is part of the asynchronous mirroring setup.
- If the logical volume type is `aio_cache` and if the logical volume is mirrored, you must use the passive mirror write consistency policy.

- The rootvg volume group cannot be configured for asynchronous mirroring.

To set up asynchronous mirroring for a mirror pool, follow these high level steps:

1. Create a scalable volume group with super strict mirror pools and add local disks.
2. Add all local disks into mirror pool
3. Add RPV devices to the mirror pool
4. Add remote mirror copy in the volume group
5. Add a logical volume of type aio_cache to a local mirror pool. The logical volume can have two copies if there two local mirror pools. The logical volume cannot have a remote copy.
6. Set up asynchronous mirroring for the mirror pool using the IO cache logical volume from the other mirror pool. Two peer mirror pools must reside at different sites.

Configuring RPV client and server pairs for a standard volume group at a node on each site

To configure a standard geographically mirrored volume using the GLVM functions, you must configure an RPV server site name, RPV server, RPV client, and add RPVs to the volume group.

To configure a standard geographically mirrored volume using the GLVM functions, complete the following tasks:

1. Configure an RPV server site name.

On all nodes at the remote site, configure an RPV Server site name.

A *site name* is an attribute that is assigned to all RPV servers configured on the remote site. In a PowerHA® SystemMirror® cluster, the RPV server site name must match the site name in PowerHA® SystemMirror®.

2. Configure an RPV server.

On one node at the remote site, configure an RPV server for the physical volumes you plan to configure as geographically mirrored volumes.

Configuring an RPV server makes the physical volume available for RPV client configuration on the node on the local site.

When you configure an RPV server, you select a physical volume with its PVID from the list of available disks on this site (that is, disks that are not assigned to a volume group), and specify the service IP addresses of all RPV clients on the TCP/IP network through which the LVM from the local site should access the disk on the node at the remote site.

You also specify that the RPV server should start immediately, which indicates that the RPV server is in the *available state*. The field **Configure Automatically at System Restart?** should be set to **No**.

3. Configure an RPV client.

On the local site, configure an RPV client for each RPV server. Configuring an RPV client for the previously configured RPV server establishes the physical volume as an RPV.

When you configure an RPV client, you first specify the service IP addresses that are to be used for RPV client-server communications. Then you select a remote physical volume (and its PVID) for which you are configuring an RPV client. Once you select or type in these values, you proceed to add a new RPV client for the previously configured RPV server.

You also specify that the RPV client should start immediately, which indicates that the RPV client is in the *available state*.

You can perform this tasks by using the **Devices** SMIT interface (the SMIT fast path is `smit devices`).

4. Add RPVs to the volume group.

On one participating node at the local site, add already configured RPVs to the volume group.

When at a later stage you add a mirror copy, the volume group becomes a geographically mirrored volume group. Import the volume group on other nodes.

You perform this task by using the **GLVM Utilities** SMIT interface.

Extending the geographically mirrored standard volume group to other nodes in the cluster

After you have completed the high-level steps 1- 4 in the section (Client-Server Pairs for a Standard Volume Group at a Node on Each Site), extend the configuration to other nodes in the cluster.

To extend the standard geographically mirrored volume group to other nodes in the cluster:

1. Configure an RPV server (in the defined state) on each node at the remote site that might end up activating the volume group due to a fallover.
2. Configure an RPV client (in the defined state) on each node at the local site that might end up activating the volume group due to a fallover.
3. To make the mirroring function work both ways in a PowerHA® SystemMirror® cluster, configure an RPV client on each node at the remote site. Configure an RPV server on each node at the local site.
4. Vary off the volume groups and update the volume groups definitions on the nodes at the other site by running the **importvg** command.

Result

The example of this procedure is included in the section Sample procedure: Preparing a volume group for inclusion in the PowerHA® SystemMirror® resource group.

Configuring RPV client-server pairs for an enhanced concurrent volume group on a node at each site

If you have configured a PowerHA® SystemMirror® cluster that spans two sites, follow this procedure to configure an enhanced concurrent mode volume group that is also geographically mirrored (that is, has RPVs added to it).

To configure a geographically mirrored enhanced concurrent volume group:

1. While the cluster services are running, configure an enhanced concurrent volume group on the disks at the local site, using the AIX® LVM utilities.
2. Using the GLVM Utilities in SMIT, on the node at the local site, add RPVs to the volume group. This allows the disks at the remote site to be accessed from the local site. This task entails the following steps:
 - a. On the local site, add the RPV clients to the volume group. For more information, see [“Configuring remote physical volume servers and clients” on page 52](#).
 - b. On the local site, use the GLVM Utilities to add a volume group mirror copy to the RPVs. For more information, see [“Adding a remote site mirror copy to a volume group” on page 57](#).
 - c. Import the volume group to all the remaining nodes at the local site.
 - d. On the local site, vary off the volume group, and configure all the RPV clients into the defined state.
 - e. On the remote site, configure all the RPV servers into the defined state.
3. Using the GLVM Utilities in SMIT, on the node at the remote site, add RPVs to allow the disks at the local site to be accessed from the remote site. For more information, see [“Configuring RPVs, mirror copies and geographically mirrored logical volumes” on page 56](#).
4. Import the volume group to all of the nodes at the remote site.
5. Vary off the volume group on all nodes at the remote site.
6. Add the enhanced concurrent geographically mirrored volume group to a PowerHA® SystemMirror® resource group.

Extending an existing enhanced concurrent volume group to nodes at remote site

When starting with an existing enhanced concurrent volume group (configured within just one site), extend this group to both sites.

To extend an existing enhanced concurrent volume group to the nodes at a remote site:

1. Extend the cluster to contain another site. Add a site name, and then add the nodes at the second site to the PowerHA® SystemMirror® cluster.
2. Add RPVs to the volume group to allow the disks at the remote site to be accessed from the local site. To do this:
 - a. On the local site, add the RPVs to the volume group.
 - b. On the local site, use the GLVM Utilities to add a volume group mirror copy to the RPVs.
 - c. On the local site, vary off the volume group and configure all RPV clients into the defined state.
 - d. On the remote site, configure all the RPV servers into the defined state.
3. Add RPVs to allow the disks at the local site to be accessed from the remote site.
4. Import the volume group to all of the nodes at the remote site.
5. Vary off the volume group on all nodes at the remote site.
6. Add the enhanced concurrent geographically mirrored volume group to a PowerHA® SystemMirror® resource group.

Configuring asynchronous mirroring

Follow these steps to set up asynchronous mirroring for a mirror pool.

Before configuring asynchronous mirroring, read the information in the section Configuration overview for asynchronous mirroring.

These instructions assume that you are starting with an ordinary volume group and changing it into an asynchronously mirrored GLVM volume group. You can also configure asynchronous mirroring of an existing GLVM volume group. In the latter case, you would only need to perform a subset of these steps, which are summarized at the end of this section.

1. Asynchronous mirroring requires the volume group to be in Scalable VG format. If your volume group is in Ordinary or Big VG format, then you must convert it to Scalable VG format before you can configure it for asynchronous mirroring. You can change your volume group to Scalable VG format by using the **chvg** command:

```
chvg -G datavg
```

Note that the volume group must be varied offline before running this command, and there are other considerations that are described in the **chvg** man page. Refer to the AIX® documentation for further information about converting your existing volume groups to Scalable VG format.

2. If you have not already done so, prevent the volume group from being varied online automatically during system startup by using the **chvg** command:

```
chvg -a n datavg
```

Refer to the AIX® documentation for further information about preventing volume groups from being activated during system startup.

3. Asynchronous mirroring requires the bad block relocation policy to be turned off for the volume group. You can turn off bad block relocation by using the **chvg** command:

```
chvg -b n datavg
```

Refer to the AIX® documentation for further information about the bad block relocation policy and how to change it.

4. Asynchronous mirroring requires the volume group to be configured to use super strict mirror pools. You can configure super strict mirror pools by using the **chvg** command:

```
chvg -M s datavg
```

5. Asynchronous mirroring also requires the local and remote disks to belong to separate mirror pools. At this point, you should define a mirror pool for the local disks at the production site. You need to choose a

name for the mirror pool. This name only needs to be unique within the volume group, so you may choose to use the same name for more than one volume group. For example, you may want to use the site name as part of the mirror pool name, to make it more meaningful. You can define a mirror pool for the local disks by using the **chpv** command:

```
chpv -p Poughkeepsie hdisk10 hdisk11
```

6. Configure your existing logical volumes to belong to the mirror pool that you just created. You also need to turn off bad block relocation for each logical volume. You can accomplish both things by using the **chlv** command.

```
chlv -m copy1=Poughkeepsie -b n dataologlv
chlv -m copy1=Poughkeepsie -b n datafslv
```

Refer to the AIX® documentation for further information about using super strict mirror pools.

7. At this point, you should have an ordinary volume group that is configured for super strict mirror pools, and all of the local disks at the production site belong to a mirror pool named "Poughkeepsie." You can confirm this by using the **lsmv** command.

```
lsmv -A datavg
VOLUME GROUP:      datavg          Mirror Pool Super Strict: yes
MIRROR POOL:      Poughkeepsie    Mirroring Mode:          SYNC
```

8. Define the RPV clients and RPV servers that are required for your cluster. On each node, you define an RPV server for each local disk that belongs to the volume group and you define an RPV client for each remote disk that belongs to the volume group. Then, the RPV client-server pairs that you create enable LVM to access remote physical volumes as if they were ordinary local disks. Refer to the section Configuring remote physical volume servers and clients for further details.
9. After you have created RPV clients and servers, you must add the remote physical volumes, which are identified by their RPV client names, to the volume group. You can define the mirror pool for the remote physical volumes in the same step. The instructions for doing this are described in the section Adding remote physical volumes to a mirror pool. This step uses the **extendvg** command to add remote physical volumes and place them into their own mirror pool at the same time. If you prefer, you may run the **extendvg** command from the command line.

```
extendvg -p Austin datavg hdisk31 hdisk32
```

In this example, the remote physical volumes hdisk31 and hdisk32 belong to a mirror pool named "Austin." Refer to the AIX® documentation for further information about using the **extendvg** command.

10. After you add the remote physical volumes to your volume group, you must add mirror copies of your logical volumes to them. The instructions for doing this are described in the section Adding a remote site mirror copy to a mirror pool. If you prefer, run the **mirrorvg** command from the command line.

```
mirrorvg -c 2 -p copy2=Austin datavg
```

In this example, a second mirror copy of each logical volume is created on the disks which reside in the Austin mirror pool. Refer to the AIX® documentation for further information about using the **mirrorvg** command.

11. Asynchronous mirroring requires a logical volume of type aio_cache to serve as the cache device. Create this logical volume using the usual steps to create a logical volume, except specify aio_cache as the logical volume type and the logical volume must reside on disks in the opposite site's mirror pool. You can perform this step using the **mklv** command.

```
mklv -y datacache1v1 -t aio_cache -p copy1=Poughkeepsie -b n -w p datavg 100
```

In this example, the cache logical volume resides in the Poughkeepsie mirror pool. During asynchronous mirroring, it is used for caching data to the disks in the Austin mirror pool. If the logical volume type is aio_cache and if the logical volume is mirrored, you must use the passive mirror write consistency policy.

12. Configure a mirror pool to use asynchronous mirroring by following the instructions that are described in the section Convert to asynchronous mirroring for a mirror pool. If you prefer, run the **chmp** command from the command line.

```
chmp -A -m Austin datavg
```

In this example, the Austin mirror pool is configured to use asynchronous mirroring. The **chmp** command automatically determines that the `datacache1v1` logical volume is the cache device to use, because it resides in the opposite site's mirror pool.

13. You can also choose to configure asynchronous mirroring for the Poughkeepsie mirror pool. To accomplish this, create a logical volume in the Austin mirror pool to serve as the cache device.

```
mk1v -y datacache1v2 -t aio_cache -p copy1=Austin -b n -w p datavg 100
```

14. Configure the Poughkeepsie mirror pool to use asynchronous mirroring by using the same procedure that you followed for the Austin mirror pool.

```
chmp -A -m Poughkeepsie datavg
```

15. You can confirm that asynchronous mirroring is configured by running the **lsmp** command again.

```
lsmp -A datavg
VOLUME GROUP:      datavg           Mirror Pool Super Strict: yes

MIRROR POOL:      Poughkeepsie   Mirroring Mode:          ASYNC
ASYNC MIRROR STATE: inactive      ASYNC CACHE LV:         datacache1v2
ASYNC CACHE VALID: yes             ASYNC CACHE EMPTY:     yes
ASYNC CACHE HWM:  100              ASYNC DATA DIVERGED:   no

MIRROR POOL:      Austin           Mirroring Mode:          ASYNC
ASYNC MIRROR STATE: active         ASYNC CACHE LV:         datacache1v1
ASYNC CACHE VALID: yes             ASYNC CACHE EMPTY:     no
ASYNC CACHE HWM:  100              ASYNC DATA DIVERGED:   no
```

Result

Now when the volume group is varied online at the Poughkeepsie site, the local disks in the Poughkeepsie mirror pool are updated synchronously, just as if this were an ordinary volume group, and the remote disks in the Austin mirror pool are updated asynchronously, using the cache device on the local Poughkeepsie disks. Likewise, when the volume group is varied online at the Austin site, the local disks in the Austin mirror pool are updated synchronously, and the remote disks in the Poughkeepsie mirror pool are updated asynchronously, using the cache device on the local Austin disks.

If you already have an existing GLVM volume group that is configured for synchronous mirroring, then you may decide to reconfigure it for asynchronous mirroring, too. First, you must follow the required steps that are described earlier to convert the volume group to Scalable VG format, turn off bad block relocation, configure the volume group to use super strict mirror pools, configure a mirror pool of disks at each site, and configure the logical volumes to belong to the new mirror pool with bad block relocation turned off.

Next, you must create a logical volume of type `aio_cache` for each mirror pool that will be configured for asynchronous mirroring. Finally, you must follow the instructions that are described in the section Convert to asynchronous mirroring for a mirror pool, or run the **chmp** command from the command line, to configure either or both of the mirror pools to use asynchronous mirroring.

Configuring remote physical volume servers and clients

These topics discuss configuring remote physical volume servers and clients.

The RPV device driver is the software layer after LVM that is used by LVM to maintain geographically remote copies of the data. Similar to standard, or local physical volumes (PVs), remote physical volumes (RPVs) can become part of a volume group, a logical volume, and a mirror copy of a logical volume. RPVs differ from local physical volumes in that they have a site name associated with them. This is the site that hosts the RPV server.

Using SMIT you can:

- Configure RPV server and client pairs for volume groups
- Display the remote physical volumes (RPVs) that are configured in the cluster, and

- List the details of the RPVs that are used.
- You can also change the state of the RPV server from defined (configured to the system but not started) to available by specifying in SMIT that the previously configured RPV server should start immediately, which indicates that the RPV server is in the available state.

Use the SMIT menu, **Devices**, to configure the RPV clients and servers. Before creating RPVs, configure volume groups and logical volumes using the standard AIX® LVM SMIT panels.

Configuring remote physical volume servers

To configure a remote physical volume server, you need to configure a volume group, logical volume, site name, and finally the server itself.

Configuring a site name on the remote site

You must define an RPV server site name before you create an RPV server.

To configure a site name:

1. Enter `smit rpvserver`
2. In SMIT, select **Remote Physical Volume Server Site Name Configuration > Define / Change / Show Remote Physical Volume Server Site Name** and press Enter.
3. Enter a field value as follows:

| <i>emote Physical Volume Server Site Name field</i> | |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Remote Physical Volume Server Site Name | Enter the name for the RPV server site. Use no more than 64 alphanumeric characters or underscores; do not use a leading numeric. Do not use reserved words. If you have previously configured sites in PowerHA® SystemMirror®, the name of the site for this node displays in the field. |

Configuring an RPV server

By configuring RPV servers, you make physical volumes available to remote RPV clients. You can configure and manage RPV servers for several physical volumes in the same operation.

To configure an RPV server:

1. Enter `smit rpvserver`
2. In SMIT, select **Add Remote Physical Volume Servers** and press Enter.
A picklist of all configured and available physical disks appears. These are the disks that are not assigned to a volume group. For example:

| <i>Available physical disks</i> | |
|---------------------------------|-----------------------------------|
| Physical Volume | Physical Volume Identifier |
| hdisk3 | 00010006859ece05 |
| hdisk4 | 0001000181810cff |
| hdisk5 | 0001000685a10b38 |
| hdisk6 | 0001000685a23307 |
| hdisk7 | 0001000685a35768 |

3. Select one or more physical disks for which you want to configure an RPV server and press Enter.
4. Enter field values as follows:

| <i>RPV fields</i> | |
|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Physical Volume Identifiers | The list of selected PVIDs displays here. An RPV server device is created for each physical volume. |
| Remote Physical Volume Client Internet Address | This is a service IP address from which a connection to the RPV client is allowed. Type in one or more addresses in this field, or press F4 to select from the list of known IP addresses defined on the remote node (that is, on the node on which the RPV client is configured). |
| Configure Automatically at System Restart? | <p>No is the default. For GLVM for PowerHA® SystemMirror® Enterprise Edition, set this field to No.</p> <ul style="list-style-type: none"> If you select No, the system does not automatically configure the RPV server at system restart. If you select Yes, the system automatically brings online the RPV server on the node at system restart. This puts the RPV server into the <i>available</i> state at system restart. |
| Start New Devices Immediately? | <p>Yes is the default.</p> <ul style="list-style-type: none"> If you select Yes, the newly created RPV server devices are configured into the <i>available</i> state immediately. If you select No, the RPV server devices remain in the <i>defined</i> state. |

5. Press Enter. The system configures an RPV server and displays the RPV server name and its current state: *defined* or *available*. Run the **lspv** command to view the RPVs that you have configured.

Result

Note: You can change the state of the RPV server later, using the **Configure Defined Remote Physical Volume Servers** panel in SMIT.

Configuring remote physical volume clients

By configuring an RPV client, you create what appears to the system as a physical volume device on the local site. When the LVM tries to access this device, the RPV client redirects the I/O to the RPV server on the remote site.

You can configure and manage several RPV clients for several physical volumes in the same operation.

Note: Before configuring an RPV client, ensure that the corresponding RPV server is in the available state on one of the nodes at the other site, while it is in the defined state on all other nodes at that site.

To configure an RPV client:

1. Enter `smit rpvclient`
2. In SMIT, select **Add Remote Physical Volume Clients** and press Enter.
3. Enter field values as follows:

| <i>Add Remote Physical Volume Clients</i> | |
|------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Remote Physical Volume Server Internet Address | Type in up to four comma-separated IP addresses in this field or select addresses from the picklist. This is an IP address on the node on the remote site on which the RPV server is configured. |

| Field | Value |
|-------|--------------------------------------------------------------------------------------|
| | This address enables a communication path between the RPV client and the RPV server. |

- Press Enter. The system connects to the RPV server and displays a list of local IP addresses from which some or all of the RPV servers on the node at the remote site will accept connection requests from the RPV clients. Select the local IP address from the list and press Enter.

You can select more than one service IP address from the list since GLVM for PowerHA® SystemMirror® Enterprise Edition supports up to four XD_data mirroring networks.

The system displays a list of physical volume names that are available and known on the node on the remote site. This is the node that hosts the RPV server address selected in the previous step. The list of physical volumes is similar to the following example:

```

These remote physical volumes are available at site somerville:
(the physical volume names are as they are known on the host named davis)

Physical Volume      Physical Volume Identifier#
-----
hdisk10             0001000181810cff0000000000000000
hdisk11             00010006859ece050000000000000000
hdisk12             0001000685a10b380000000000000000

```

- Select one or more physical volumes from the list and press Enter. The **Add Remote Physical Volume Clients** SMIT screen displays.
- Enter field values as follows:

| <i>Add Remote Physical Volume Clients fields</i> | |
|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Remote Physical Volume Server Internet Address | This field is filled in with the service IP addresses on the node on the remote site on which the RPV server is configured. The IP addresses belong to the RPV server containing the physical volume; they enable communication paths between the RPV client and the RPV server. |
| Remote Physical Volume Local Internet Address | This field is filled in with the local IP addresses of the RPV clients that you are configuring on the local node. These are the local IP address(es) from which the system connects to the RPV server. The RPV server accepts only connections from its configured RPV client IP addresses. |
| Physical Volume Identifiers | The list of selected physical volumes displays here. An RPV client device is created for each of the selected physical volumes that reside on the node on which the RPV server is configured on the remote site. |
| I/O Timeout Interval (Seconds) | How many seconds the RPV client should wait for the RPV server to respond to an I/O request before giving up and assuming that the RPV server is down on the node or is unreachable. The default is 180 seconds. In GLVM for PowerHA® SystemMirror® Enterprise Edition, this time automatically defaults to the Total Event Duration Time setting that you can specify for the config_too_long warning timer in the Change/Show Time Until Warning panel in SMIT. |
| Start New Devices Immediately? | Yes is the default. If you select Yes , the newly created RPV client devices are configured into the available state immediately. If |

| Field | Value |
|-------|--------------------------------------------------------------------------------------|
| | you select No , the RPV client devices are configured into the defined state. |

Using this SMIT panel, you can add an RPV client that is in the available state, or the client that is in the defined state. You can change the state of the RPV client later, using the **Configure Defined Remote Physical Volume Clients** panel in SMIT.

- Press Enter. The system displays the names of the newly created RPV clients. Or, run the **lspv** command to view the RPVs that you have configured.

Result

You have configured RPV clients for one or more physical volumes that reside on the node on which the RPV server is configured on the remote site.

Configuring RPVs, mirror copies and geographically mirrored logical volumes

With *Geographic Logical Volume Manager (GLVM)* utilities you can configure, view, change, and verify geographic mirrors of volume groups and logical volumes on geographically remote sites.

Using the **GLVM Utilities** SMIT panels, you can create a mirror copy of the logical volume or a copy of the volume group that resides entirely on a specified geographically remote site. You can also change and verify geographically mirrored volume groups and logical volumes.

Note: The AIX® Storage Management component provides several ways to configure disk mirrors: the **GLVM Utilities** SMIT and the standard AIX® LVM SMIT (or commands). While it is possible to use the standard LVM commands and SMIT panels to correctly add mirror copies, using the **GLVM Utilities** SMIT interface prevents you from configuring the mirror copies in such a way that data would be lost if any one of the physical volumes failed. In the **GLVM Utilities** SMIT, it is easier to make valid selections from the picklists and the system verifies your selections to ensure that the geographic mirroring function is configured correctly. The **GLVM Utilities** SMIT helps you to avoid mistakes that would keep a complete copy of the data from being available after a site failure.

Configuration prerequisites for GLVM utilities

Before configuring volume groups or logical volumes as geographically mirrored, ensure that you have configured:

- Volume groups on the nodes on the local site. Each volume group could either be a standard volume group or an enhanced concurrent volume group. It should already contain one or more logical volumes established with local physical volumes. You configure the volume groups, logical volumes and add local copies using the standard AIX® LVM commands or SMIT panels. Also, it is a good idea to configure file systems and to have application data included in the volume groups.
- Logical volumes created with the Super Strict inter-disk allocation policy. See the AIX® documentation for information on changing the inter-disk allocation policy.
- Physical volumes and remote physical volumes (RPVs).
- One local and only one remote site.

Using GLVM utilities: Managing RPVs and remote site mirror copies

Once you have configured volume groups and logical volumes, the SMIT fields in **Geographically Mirrored Volume Groups** SMIT panels allows you to perform several different tasks.

The LVM allows you to create up to three copies (the original copy and two additional copies); copies can be added at the same time and must reside on two sites: up to two copies on the local site and one on the remote site.

Note: You can use the `smit glvm_utils` fastpath for accessing the **GLVM Utilities** SMIT. This fastpath is valid for all the procedures that follow.

Adding remote physical volumes to a volume group

To configure a geographically mirrored volume group, you add RPVs to it.

Follow these steps:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Volume Groups > Add Remote Physical Volumes to a Volume Group** and press Enter.
3. Select a volume group name from the picklist and press Enter. SMIT displays a list of available sites on which RPVs are configured. (Note: Since PowerHA® SystemMirror® Enterprise Edition for GLVM allows the configuration of only one remote site, only one site displays in this list).
4. Select a site name and press Enter. SMIT displays a list of all free remote physical volumes configured on the selected site, as in the following example:

```
#All Free Physical Volumes in site chicago:
# RPV   PVID                               Site
hdisk6  0001234567891122                       chicago
hdisk3  0009876543213367                       chicago
hdisk9  0005678943217896                       chicago
```

5. Select one or more remote physical volumes you want to add to the volume group and press Enter.
6. Enter field values as follows:

| <i>Add Remote Physical Volumes to a Volume Group fields</i> | |
|-------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Volume Group Name | This is the name of the volume group to which you want to add RPVs to make it a geographically mirrored volume group. |
| Force | The default is No . If you select Yes , the system forces the remote physical volume to be added to the specified volume group unless it is a member of another volume group in the Device Configuration Database or of a volume group that is active. |
| Remote Physical Volume Name | The name of the remote physical volume, such as <code>hdisk3</code> . More than one name can be selected in this field. Remote physical volume names are typically in the form of <code>hdiskx</code> where <code>x</code> is a system-wide unique number. This name is assigned when the RPV client is created. Remote physical volumes have a site attribute. The remote <code>hdisk</code> number can differ from the local physical volume number, but the underlying PVID is always the same. |

7. Press Enter. The system runs the `extendvg` command and adds the RPV names to the list of physical volumes configured for this volume group. The RPVs can be used to create a geographically remote mirror copy of one or all logical volumes in the volume group.

Adding a remote site mirror copy to a volume group

To make a remote mirror copy of all logical volumes in a volume group on physical volumes at a remote site, you add a mirror copy to a volume group.

The GLVM utilities allow you to add multiple mirror copies, and ensure that they are created at the remote site, whereas the standard AIX® SMIT for the LVM lets you create mirror copies on either local or remote site.

To add a remotely located mirror copy to a volume group:

1. Enter `smit`.
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Volume Groups > Add a Remote Site Mirror Copy to a Volume Group** and press Enter.
SMIT displays all non-root standard and enhanced concurrent volume groups that have remote physical volumes (RPVs) assigned to them.
3. Select a volume group for which you want to add a mirror copy at a remote site and press Enter. SMIT displays a list of defined sites.
4. Select a site on which a mirror copy should reside and press Enter. SMIT displays a list of RPVs defined on this site.
5. Select one or more unassigned RPVs. Ensure that the number of RPVs you select is sufficient for creating a mirror copy. Press Enter.
6. Enter field values as follows:

| <i>RPV fields</i> | |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Volume Group Name | The volume group name for which you want to create a mirror copy at the remote site. |
| Remote Physical Volume Names | The list of selected RPVs that reside on the nodes at the remote site on which the Geographic Logical Volume Manager will create a mirror copy of the volume group. |
| Mirror Sync Mode | <p>The default is Foreground.</p> <p>After you create the mirror copy, the system synchronizes the mirrors by copying data from the original copy to the newly created copy. Select the option for the synchronization of mirrors:</p> <ul style="list-style-type: none"> • Foreground. The system synchronizes the mirrors immediately and all other processes wait. • Background. The system synchronizes the mirrors immediately but other processes can continue in the background. • No Sync. The system does not synchronize the mirrors at this time. |
| Keep Quorum Checking On? | <p>The default is No. Indicate whether you want the system to ensure data consistency between mirror copies of a logical volume during normal I/O processing:</p> <div style="border: 1px solid blue; padding: 5px; margin: 10px 0;"> <p>For GLVM for PowerHA® SystemMirror® Enterprise Edition, the quorum should not be set.</p> </div> <ul style="list-style-type: none"> • Yes. The system continues to check for quorum while attempting to create a mirror copy. • No. The system does not check for the disk quorum. |

7. Press Enter. The GLVM verifies that all logical volumes in the volume group are mirrored with the Super Strict inter-disk allocation policy, and makes a mirror of all logical volumes in the volume group using the RPVs located on the specified remote site. The GLVM ensures that the complete mirror copy is created on the remote site.

The SMIT panel fails and issues an error if one of the logical volumes in the volume group does not have the Super Strict inter-disk allocation policy.

- Now you can verify mirror copy site locations for a volume group.

Removing a remote site mirror copy from a volume group

With this SMIT option, you remove one copy of a geographically mirrored logical volume which is located on one set of local physical volumes or RPVs. Replacing a physical volume or an RPV can be performed with the standard AIX® `replacepv` command or using the similar options in the Geographically Mirrored Volume Groups SMIT panel.

Follow these steps:

- Enter `smit`
- In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Volume Groups > Remove a Remote Site Mirror Copy from a Volume Group** and press Enter.
SMIT displays all non-root standard and enhanced concurrent volume groups that have RPVs assigned to them and that have mirrors created at a remote site.
- Select a volume group for which you want to remove a mirror copy at a remote site and press Enter. SMIT displays a list of mirror copies.
- Select a mirror copy which you want to remove and press Enter
- Enter field values as follows:

| <i>Remove a Remote Site Mirror Copy from a Volume Group fields</i> | |
|--------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Volume Group Name | The volume group name for which you want to remove a mirror copy at the remote site. |
| Remote Physical Volume Names | The list of selected RPVs that reside on the nodes at the remote site for which the Geographic Logical Volume Manager will remove a mirror copy of the volume group. |
| Number of Copies of each Logical Partition | The default value for the number of copies for each logical volume is 1 . Specify the number of copies of each logical partition that you want to have in the volume group. For example, if three copies exist, and you specify 2 in this field, one copy will be removed. If two copies exist, none will be removed. If one copy exists, it will not be removed. |

- Press Enter. The SMIT asks "**Are you sure?**" The GLVM removes one copy at a time. If a mirror copy of the logical volume spans more than one disk, the portion of the logical volume copy that resides on a disk not listed by you is also removed.

Verifying mirror copy site locations for a volume group

Once you have finished configuring mirror copies, you should verify that they are properly configured.

To verify that mirror copies for geographically mirrored volume groups are properly configured:

- Enter `smit`
- In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Volume Groups > Verify Geographic Volume Group Information** and press Enter.
SMIT displays a list of volume groups.
- Select the volume groups from the list, or enter the names of the volume groups for which you want to verify geographic mirrors and press Enter.

Result

The verification command verifies that the logical volumes are configured with the Super Strict inter-disk allocation policy.

Managing geographically mirrored volume groups with mirror pools

These topics describe the maintenance tasks for geographically mirrored volume groups with mirror pools.

Adding remote physical volumes to a mirror pool

With this SMIT option, you add remote physical volumes to an already existing mirror pool.

To add a remote physical volume to a mirror pool:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Manage Geographically Mirrored Volume Groups with Mirror Pools > Add Remote Physical Volumes to a Mirror Pool** and press Enter.
SMIT displays a list of all configured volume groups that have available RPVs.
3. Select a Volume Group.
4. Select one or more remote physical volumes.
5. Enter the field values as follows:

| <i>Add Remote Physical Volumes to a Mirror Pool fields</i> | |
|------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| VOLUME GROUP Name | The name of the volume group that contains the mirror pool. |
| Force | Possible values of Yes or No . Yes forces the physical volume to be added to the group. |
| REMOTE PHYSICAL VOLUME names | The name of the remote physical volume that you are adding. |
| MIRROR POOL Name | The name of the mirror pool. If you want to add a remote physical volume to a new mirror pool, enter the name of the new mirror pool. |

You must enter a mirror pool name.

6. Press Enter.

Removing remote physical volumes from a mirror pool

Use this task to remove a remote physical volume from a mirror pool.

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Manage Geographically Mirrored Volume Groups with Mirror Pools > Remove Remote Physical Volumes from a Mirror Pool** and press Enter.
SMIT displays a list of all configured volume groups.
3. Select a volume group.
4. Select a mirror pool.
5. Select one or more remote physical volumes.
6. Enter the field values as follows:

| <i>Remove Remote Physical Volumes from a Mirror Pool fields</i> | |
|-----------------------------------------------------------------|-------------------------------------------------------------|
| Field | Value |
| VOLUME GROUP Name | The name of the volume group that contains the mirror pool. |

| Field | Value |
|---------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| FORCE deallocation of all partitions on this physical Volume? | Deallocates the existing logical volume partitions and then deletes resultant empty logical volumes from the specified physical volumes. |
| REMOTE PHYSICAL VOLUME Names | Remote physical volumes that you are removing. |

7. Press **Enter**.

Adding a remote site mirror copy to a mirror pool

With this SMIT option, you create a mirror copy of an individual logical volume to a mirror pool.

To add a remotely located mirror copy to a mirror pool:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Manage Geographically Mirrored Volume Groups with Mirror Pools > Add a Remote Site Mirror Copy to a Mirror Pool** and press **Enter**.
SMIT displays a list of all configured volume groups that have available RPVs.
3. Select a Volume Group.
4. Select a Mirror Pool.
5. Enter the field values as follows:

| Field name | Description |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VOLUME GROUP Name | The name of the volume group that contains the mirror pool. |
| Foreground, Background, No Sync | Specifies whether synchronization of the volume group should occur immediately (Foreground), in the Background, or not at all. As portions of the mirrors become synchronized, they are available for immediate use. If you choose not to synch, the mirror copy may exist for a logical volume, but is not used until it has been synchronized. |
| MIRROR POOL Name | The name of the mirror pool that you want to convert to asynchronous mirroring. |
| Number of COPIES of each logical partition | Specifies the minimum number of copies that each logical volume must have. Minimum value is 2 and 3 is the maximum value. A value of 1 is ignored. |
| Keep Quorum Checking On? | By default, when a volume group's contents becomes mirrored, volume group quorum is disabled. To keep the volume group quorum requirement after mirroring is complete, select Yes. |

6. Press **Enter**.

Removing a remote site mirror copy from a mirror pool

With this SMIT option, you remove one copy of a geographically mirrored logical volume which is located in a mirror pool.

To remove a remote physical volume from a mirror pool:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Manage Geographically Mirrored Volume Groups with Mirror Pools > Remove a Remote Site Mirror Copy from a Mirror Pool** and press **Enter**.
SMIT displays a list of all configured volume groups.

3. Select a volume group for which you want to remove a mirror copy at a remote site and press Enter. SMIT displays a list of mirror copies.
4. Select a mirror copy which you want to remove and press Enter
5. Enter field values as follows:

| <i>Remove a Remote Site Mirror Copy from a Mirror Pool</i> | |
|------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Volume Group Name | The volume group name for which you want to remove a mirror copy in the mirror pool. |
| Remote Physical Volume Names | The list of selected RPVs that reside on the nodes in the mirror pool for which the Geographic Logical Volume Manager will remove a mirror copy of the volume group. |
| Number of Copies of each Logical Partition | The default value for the number of copies for each logical volume is 1 . Specify the number of copies of each logical partition that you want to have in the volume group. For example, if three copies exist, and you specify 2 in this field, one copy will be removed. If two copies exist, none will be removed. If one copy exists, it will not be removed. |

6. Press Enter. The SMIT asks "Are you sure?" The GLVM removes one copy at a time.

List asynchronous mirroring information for a mirror pool

You can use SMIT to display the asynchronous attributes for all volume groups and mirror pools.

To display the asynchronous attributes:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Manage Geographically Mirrored Volume Groups with Mirror Pools > Configure Mirroring Properties of a Mirror Pool > List Asynchronous Mirroring Information for Mirror Pool** and press Enter.

Result

SMIT displays a list of asynchronous attributes for all volume groups and mirror pools in this format:

| | | | |
|---------------------|----------|---------------------------|------------|
| VOLUME GROUP: | glvm_vg | Mirror Pool Super Strict: | yes |
| MIRROR POOL: | mp_bvr | Mirroring Mode: | ASYNC |
| ASYNC MIRROR STATE: | inactive | ASYNC CACHE LV: | mp_pok_lv |
| ASYNC CACHE VALID: | yes | ASYNC CACHE EMPTY: | yes |
| ASYNC CACHE HWM: | 60 | ASYNC DATA DIVERGED: | no |
| MIRROR POOL: | mp_pok | Mirroring Mode: | ASYNC |
| ASYNC MIRROR STATE: | active | ASYNC CACHE LV: | bvr_pok_lv |
| ASYNC CACHE VALID: | yes | ASYNC CACHE EMPTY: | no |
| ASYNC CACHE HWM: | 90 | ASYNC DATA DIVERGED: | no |

Convert to asynchronous mirroring for a mirror pool

You can convert a mirror pool from synchronous mirroring to asynchronous mirroring.

To convert a mirror pool:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Manage Geographically Mirrored Volume Groups with Mirror Pools > Configure Mirroring Properties of a Mirror Pool > Convert to Asynchronous Mirroring for a Mirror Pool** and press Enter.
3. Select a Volume Group.
4. Select a Mirror Pool.

- Enter the field values as follows:

| <i>Convert to Asynchronous Mirroring for a Mirror Pool fields</i> | |
|-------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| VOLUME GROUP Name | The name of the volume group that contains the mirror pool. |
| MIRROR POOL Name | The name of the mirror pool that you want to convert to asynchronous mirroring. |
| Logical Volume for IO Cache | Specifies the name of an asynchronous I/O-cache logical volume. The logical volume must be of the aio_cache type and must not reside in the mirror pool. If you leave this field blank, GLVM will attempt to find it automatically. |
| IO Cache High Water Mark Value | Specifies the I/O-cache high watermark. The value is the percent of I/O cache size. The default value is 100%. |

- Press **Enter**.

Changing asynchronous mirroring attributes for a mirror pool

You can change mirroring attributes such as the high water mark or the IO cache for a mirror pool.

To change asynchronous mirroring attributes:

- Enter `smit`
- In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Manage Geographically Mirrored Volume Groups with Mirror Pools > Configure Mirroring Properties of a Mirror Pool > Change Asynchronous Mirroring Attributes for a Mirror Pool** and press **Enter**.
SMIT displays a list of all configured volume groups.
- Select a Volume Group.
- Select a Mirror Pool.
- Enter the field values as follows:

| <i>Mirror Pool fields</i> | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------|
| Field | Value |
| VOLUME GROUP Name | The name of the volume group that contains the mirror pool. |
| MIRROR POOL Name | The name of the mirror pool that you want to change attributes for asynchronous mirroring. |
| IO Cache High Water Mark Value | Specifies the I/O-cache high watermark. The value is the percent of I/O cache size. The default value is 100%. |
| New Logical Volume for IO Cache | Displays the aio_cache type logical volumes from the same mirror pool. |

- Press **Enter**.

Converting to synchronous mirroring for a mirror pool

You can change a mirror pool to use synchronous mirroring.

To convert a mirror pool to synchronous mirroring:

- Enter `smit`

- In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Manage Geographically Mirrored Volume Groups with Mirror Pools > Configure Mirroring Properties of a Mirror Pool > Convert to Synchronous Mirroring for a Mirror Pool** and press Enter.
SMIT displays a list of all configured volume groups.
- Select a Volume Group.
- Select a Mirror Pool.
- Confirm the field values as follows:

| <i>Mirror Pool fields</i> | |
|---------------------------|--------------------------------------------------------------------------------|
| Field | Value |
| VOLUME GROUP Name | The name of the volume group that contains the mirror pool. |
| MIRROR POOL Name | The name of the mirror pool that you want to convert to synchronous mirroring. |

- Press **Enter**.

Activating an asynchronously mirrored volume group

Use this task to activate an asynchronously mirrored volume group under normal circumstances.

To activate an asynchronously mirrored volume group:

- Enter `smit`
- In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Manage Geographically Mirrored Volume Groups with Mirror Pools > Activate an Asynchronously Mirrored Volume Group > Activate under Normal Circumstances** and press Enter.
- Enter the field values as follows:

| <i>Activate under Normal Circumstances fields</i> | |
|---------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Asynchronously Mirrored VOLUME GROUP name | The name of the volume group to activate. |
| RESYNCHRONIZE stale physical partitions? | Physical partitions marked as stale must be updated to contain the same information as valid physical partitions. This process, called resynchronization, can be done at vary-on time or can be started any time while the system is running. Until the stale partitions have been rewritten with valid data, they are not used to satisfy read requests nor are they written to on write requests. Select Yes to resynchronize the partition automatically. |
| Activate volume group in SYSTEM MANAGEMENT mode? | Makes the volume group available in System Management mode only. Logical volume commands can operate on the volume group, but no logical volumes can be opened for input or output. |
| FORCE activation of the volume group? Warning--this may cause loss of data integrity. | Allows a volume group to be made active that does not currently have a quorum of available disks. All disk that cannot be brought to an active state will be put in a removed state. At least one disk must be available for use in the volume group. |
| Synchronize Logical Volumes? | Sets the synchronization characteristics for the volume group. Yes attempts to automatically synchronize stale partitions. No prohibits automatic synchronization of stale partitions. |

4. Press **Enter**.

Activating a volume group at the remote site after a production site failure

Use this menu to activate the volume group at the disaster recovery site, after the production site has failed, when using asynchronous mirroring. PowerHA® SystemMirror® will use your default set up for activating volume groups, but you can do these steps manually.

You must indicate whether or not to allow data divergence to occur.

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Manage Geographically Mirrored Volume Groups with Mirror Pools > Activate an Asynchronously Mirrored Volume Group > Activate at Disaster Recovery Site after Production Site Failure** and press **Enter**.
3. Enter the field values as follows:

| <i>Activate at Disaster Recovery Site after Production Site Failure fields</i> | |
|---------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Asynchronously Mirrored VOLUME GROUP name | The name of the volume group. |
| Activate volume group in SYSTEM MANAGEMENT mode? | Makes the volume group available in System Management mode only. Logical volume commands can operate on the volume group, but no logical volumes can be opened for input or output. |
| FORCE activation of the volume group? Warning--this may cause loss of data integrity. | Allows a volume group to be made active that does not currently have a quorum of available disks. All disk that cannot be brought to an active state will be put in a removed state. At least one disk must be available for use in the volume group. |
| Synchronize Logical Volumes? | Sets the synchronization characteristics for the volume group. Yes attempts to automatically synchronize stale partitions. No prohibits automatic synchronization of stale partitions. |
| Allow activation with missing data updates? Warning--may lead to data divergence. | Allows data divergence. This option only takes effect when you try to bring the volume group online while the cache at the opposite site might contain nonmirrored data updates and that cache is not accessible. If the system detects that you might use back-level data and you do not specify Yes in this field, the activation will fail. |

4. Press **Enter**.

Activating a volume group after production site recovery

Use the Activate after Production Site Recovery only after data divergence has occurred and you want to use settings other than your default to activate volume groups.

After recovering the production site, you should use this menu to tell the varyonvg command which site's copy of the data should be preserved. Any non-mirrored updates which were written to the other site's copy will be backed out, meaning they will be lost, in order to merge the two diverged versions of the volume group back together.

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Manage Geographically Mirrored Volume Groups with Mirror Pools > Activate an Asynchronously Mirrored Volume Group > Activate after Production Site Recovery** and press **Enter**.
3. Enter the field values as follows:

| <i>Activate after Production Site Recovery fields</i> | |
|-----------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Asynchronously Mirrored VOLUME GROUP name | The name of the volume group. |
| RESYNCHRONIZE stale physical partitions? | Physical partitions marked as stale must be updated to contain the same information as valid physical partitions. This process, called resynchronization, can be done at vary-on time or can be started any time while the system is running. Until the stale partitions have been rewritten with valid data, they are not used to satisfy read requests nor are they written to on write requests. Select Yes to resynchronize the partition automatically. |
| Activate volume group in SYSTEM MANAGEMENT mode? | Makes the volume group available in System Management mode only. Logical volume commands can operate on the volume group, but no logical volumes can be opened for input or output. |
| FORCE activation of the volume group? Warning--this may cause loss of data integrity. | Allows a volume group to be made active that does not currently have a quorum of available disks. All disk that cannot be brought to an active state will be put in a removed state. At least one disk must be available for use in the volume group. |
| Synchronize Logical Volumes? | Sets the synchronization characteristics for the volume group. Yes attempts to automatically synchronize stale partitions. No prohibits automatic synchronization of stale partitions. |
| Recover from data divergence, if necessary? If YES, specify which copy to preserve. | Answer NO if you do not want to allow data divergence recovery to occur. However, if data divergence has occurred, answering NO causes this operation to fail. If you do want to allow data divergence recovery to occur (YES), then you must indicate whether to preserve the data copy that resides at the local or remote site. In this case, local refers to the local site, where you are running this SMIT menu, and remote refers to the opposite site. This operating will preserve the data at the site you select and backs out the non-mirrored updates at the opposite site. |
| FORCE data divergence recovery, if necessary? Warning--this may cause loss of data integrity. | Yes forces data divergence recovery, even if the data copy contains stale physical partitions. Be aware that this can cause loss of data integrity. |

4. Press **Enter**.

Adding a remote site mirror copy to a logical volume

With this SMIT option, you create a mirror copy of an individual logical volume on a set of RLVs on the remote site. The GLVM SMIT uses the `mklvcopy` command for this purpose. Typically, you do this task when you want to add a new logical volume to an already existing geographically mirrored volume group.

To add a remotely located mirror copy to a logical volume:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Add a Remote Site Mirror Copy to a Logical Volume** and press **Enter**.

SMIT displays a list of all configured logical volumes that contain no mirror copies or have one mirror copy configured, and that belong to a volume group that has available RLVs.

3. Select an appropriate logical volume and remote physical volumes for which to create a mirror of the logical volume. This increases the number of physical partitions per logical partition within the logical volume.

A logical partition may have one physical partition per logical partition, which is known as a “one copy” or “unmirrored logical volume” or it may have two or three physical partitions per logical partition, which is known as a “mirrored logical volume”.

Press Enter. SMIT displays a list of defined sites.

4. Select a site on which a mirror copy should reside and press Enter. This ensures that the mirror copy that you create will reside entirely on one site. SMIT displays a list of RLVs defined on this site.
5. Select one or more unassigned remote physical volumes (RLVs). Ensure that the number of RLVs you select is sufficient for creating a mirror copy. Press Enter.
6. Enter field values as follows:

| <i>Add a Remote Site Mirror Copy to a Logical Volume fields</i> | |
|---------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Logical Volume Name | The name of the logical volume for which you want to create a mirror copy at the remote site. |
| Remote Physical Volume Names | The list of selected RLVs that reside on the nodes at the remote site on which Geographic Logical Volume Manager creates a mirror copy of the logical volume. |
| New Total Number of Logical Partition Copies | The default is 2 . The GLVM utilities can create two copies if you specify 3 here. |
| Position on Physical Volume | The default is Middle . This option uses the default value from the standard AIX® LVM panel. |
| Range of Physical Volumes | The default is Minimum . This option uses the default value from the standard AIX® LVM panel. |
| Allocate each logical partition copy on a SEPARATE physical volume? | Super Strict . This is the required setting for GLVM for PowerHA® SystemMirror® Enterprise Edition. The Super Strict inter-disk allocation policy lets GLVM properly mirror the logical volume at the remote site. |
| SYNCHRONIZE the data in the new logical partition copies? | No . This option uses the default value from the standard AIX® LVM panel. |

7. Press Enter. The GLVM makes a mirror of the logical volumes using the RLVs located on the specified remote site.

Result

Note: If one of the logical volumes does not have the Super Strict inter-disk allocation policy, the SMIT panel issues an error.

Extending a logical volume that contains remote site mirrors

You can increase the size of a logical volume by adding unallocated physical partitions to it from within the volume group. The GLVM command for extending the logical volume uses the Super Strict inter-disk allocation policy for the logical volumes and assumes that only two sites are configured, one of which is a remote site.

To extend a logical volume that contains remote site mirrors:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Extend a Logical Volume that Contains a Remote Site Mirror** and press Enter. SMIT displays a list of logical volumes.

3. Select a logical volume to which you want to add physical partitions and press Enter.
4. Enter or review field values as follows:

| <i>Extend a Logical Volume that Contains a Remote Site Mirror fields</i> | |
|--------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Logical Volume Name | The name of the logical volume for which you want to add physical partitions. |
| Number of Additional Logical Partitions | <p>Specify the number of logical partitions you want to add to the logical volume.</p> <div style="border: 1px solid blue; padding: 5px; margin: 5px 0;"> <p>Sufficient space must exist on each mirror copy to set aside for the new physical partitions.</p> </div> <p>The system requests the additional physical partitions first from the pool of unused but available partitions configured on the physical volume and on the remote physical volume.</p> <p>If there is not enough free partitions in the pool of available partitions, the system requests them from the volume group pool.</p> |
| Position on Physical Volume | The default is Middle . This option uses the default value from the standard AIX® LVM panel. |
| Range of Physical Volumes | The default is Minimum . This option uses the default value from the standard AIX® LVM panel. |
| Allocate each logical partition copy on a SEPARATE physical volume? | Super Strict . This forces each complete copy of the logical volume to reside on separate sets of physical volumes (and remote physical volumes). |
| Synchronize the data in the new logical partition copies? | The default is No . Select No . Physical partitions marked as stale must be updated to contain the same information as valid physical partitions. This process, called resynchronization, can be done at varyon time or can be started any time while the system is running. Until the stale partitions have been rewritten with valid data, they are not used to satisfy read requests nor are they written to based on write requests. |

5. Press Enter. The system examines the logical volume and determines how to extend it by the specified number of physical partitions.
6. Verify the configuration.

Removing a remote site mirror copy from a volume group

With this SMIT option, you remove one copy of a geographically mirrored logical volume which is located on one set of local physical volumes or RPVs. Replacing a physical volume or an RPV can be performed with the standard AIX® `replacepv` command or using the similar options in the Geographically Mirrored Volume Groups SMIT panel.

Follow these steps:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Volume Groups > Remove a Remote Site Mirror Copy from a Volume Group** and press Enter.
SMIT displays all non-root standard and enhanced concurrent volume groups that have RPVs assigned to them and that have mirrors created at a remote site.

3. Select a volume group for which you want to remove a mirror copy at a remote site and press Enter. SMIT displays a list of mirror copies.
4. Select a mirror copy which you want to remove and press Enter
5. Enter field values as follows:

| <i>Remove a Remote Site Mirror Copy from a Volume Group fields</i> | |
|--------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Volume Group Name | The volume group name for which you want to remove a mirror copy at the remote site. |
| Remote Physical Volume Names | The list of selected RPVs that reside on the nodes at the remote site for which the Geographic Logical Volume Manager will remove a mirror copy of the volume group. |
| Number of Copies of each Logical Partition | The default value for the number of copies for each logical volume is 1 . Specify the number of copies of each logical partition that you want to have in the volume group. For example, if three copies exist, and you specify 2 in this field, one copy will be removed. If two copies exist, none will be removed. If one copy exists, it will not be removed. |

6. Press Enter. The SMIT asks "**Are you sure?**" The GLVM removes one copy at a time. If a mirror copy of the logical volume spans more than one disk, the portion of the logical volume copy that resides on a disk not listed by you is also removed.

Listing all geographically mirrored logical volumes

You can use SMIT to list all geographically mirrored logical volumes that are already configured.

To list all volumes:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > List Geographic Logical Volume Information** and press Enter.

Result

SMIT displays all remote physical volumes and their assignment to volume groups and logical volumes in the following format:

| # | Volume Group | Logical Volume | RPV PVID | Site |
|-----|--------------|----------------|--------------|---------|
| vg1 | lv1 | hdisk1 | 000987654321 | chicago |

Verifying mirror copy site locations for a logical volume

With this SMIT option, you verify that logical volume mirrors which have geographic copies have been configured correctly by determining that each copy has physical volumes that reside on only one site and that the logical volume has the Super Strict inter-disk allocation policy.

To verify configuration:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Verify Geographic Logical Volume Information** and press Enter.
SMIT displays a list of logical volumes.
3. Select the logical volumes from the list, or enter the names of the logical volumes for which you want to verify geographic mirrors and press Enter.
When you verify the mirror copies, the system verifies that:

- A copy of the logical volume exists on only one site: A local copy exists totally on the local site, and the remote copy totally exists on the geographically remote site.

- All logical volumes are configured with the Super Strict inter-disk allocation policy.

Result

When you select the logical volumes from the list, or enter the names of the logical volumes and press Enter, the system issues a message: "Checking logical volume <lv_name>". After a brief delay, the SMIT interface displays either a successful command completion or issues an error.

Changing the RPV servers and clients configuration

These topics describe the tasks that allow you to change the RPV configuration.

Additional configuration tasks for RPV servers

Using the `smit rpvserver fastpath`, or the SMIT path **System Management > Devices > Remote Physical Volume Servers**, you can perform several additional tasks.

Changing an RPV server site name

You must define an RPV server site name before you create an RPV server.

Note: If you want to change an already existing RPV server site name, all RPV servers on the node must be in the *defined* state. Before changing the site name, change the state of the RPV servers from available to defined. Changing the RPV server site name does not change the site name in PowerHA® SystemMirror®.

To change a site name:

1. Enter `smit rpvserver`
2. In SMIT, select **Remote Physical Volume Server Site Name Configuration > Define / Change / Show Remote Physical Volume Server Site Name** and press Enter.
3. Enter a field value as follows:

| <i>Remote Physical Volume Server Site Name fields</i> | |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Remote Physical Volume Server Site Name | Enter the name for the RPV server site. Use no more than 64 alphanumeric characters or underscores; do not use a leading numeric. Do not use reserved words. If you have previously configured sites in PowerHA® SystemMirror®, the name of the site for this node displays in the field. |

Listing all remote physical volume servers

You can list all previously configured remote physical volume servers.

To list all RPV servers:

1. Enter `smit rpvserver`
2. In SMIT, select **List All Remote Physical Volume Servers** and press Enter.

Result

The system lists all previously configured RPV servers, as in the following example:

| RPV Server | Physical Volume Identifier | Physical Volume |
|-------------------|-----------------------------------|------------------------|
| rpvserver0 | 00010006859ece05 | hdisk3 |
| rpvserver1 | 0001000181810cff | hdisk4 |

| RPV Server | Physical Volume Identifier | Physical Volume |
|------------|----------------------------|-----------------|
| rpvserver2 | 0001000685a10b38 | hdisk5 |

Changing or showing a remote physical volume server

You can change or show a remote physical volume server name and its current state.

To change or show an RPV server, complete the following steps:

1. From the command line, enter `smit rpvserver`.
2. In SMIT, select **Change/Show a Remote Physical Volume Server** and press Enter.
A list of previously configured servers displays. Select an RPV server you want to change or show and press Enter.
3. Enter field values as follows

| <i>Change/Show a Remote Physical Volume Server fields</i> | |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Remote Physical Volume Server | The name of the RPV server for which you want to change parameters. |
| Physical Volume Identifier | The Physical Volume Identifier (PVID) associated with the RPV server. This value cannot be changed. |
| Remote Physical Volume Client Internet Address | The IP address(es) of the RPV client from which a connection to the RPV server is allowed. Type in one or more new addresses in this field. The names must be separated by commas and have no spaces. Or, press F4 to select from the list of known IP addresses defined on the remote node. |
| Configure Automatically at System Restart? | For GLVM for PowerHA® SystemMirror® Enterprise Edition, set this field to No . No is the default. <ul style="list-style-type: none"> • If you select No, the system does not automatically configure the RPV server at system restart. • If you select Yes, the system automatically brings online the RPV server on the node at system restart. This puts the RPV server into the <i>available</i> state at system restart. |

4. Press Enter. The system makes changes to the RPV server and displays the RPV server name and its current state: *defined* or *available*.

Changing multiple remote physical volume servers

You can use this SMIT option to change the states of the RPV servers (*defined* and *available*), or other parameters.

To change or show multiple RPV servers:

1. Enter `smit rpvserver`
2. In SMIT, select **Change Multiple Remote Physical Volume Servers** and press Enter.
The system displays a list of previously configured RPV servers. Press F7 to select one or more RPV servers to change and press Enter.
3. Enter field values as follows:

| <i>Change Multiple Remote Physical Volume Server fields</i> | |
|-------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Remote Physical Volume Server | The names of the RPV servers for which you want to change parameters. |
| Physical Volume Identifier | The Physical Volume Identifiers (PVIDs) associated with the RPV servers. These values cannot be changed. |
| Remote Physical Volume Client Internet Address | <p>You can leave this field blank if you do not want to change it. This is the list of IP addresses from which a connection to the RPV server is allowed. Any of the addresses listed in this field can be used by each of the RPV servers.</p> <p>Type in one or more new addresses in this field. The names must be separated by commas and have no spaces. Or, press F4 to select from the list of known IP addresses defined on the remote node.</p> |
| Configure Automatically at System Restart? | <p>Do not Change is the default. For GLVM for PowerHA® SystemMirror® Enterprise Edition, select this option.</p> <ul style="list-style-type: none"> • If you select Do not Change, the RPV servers continue to behave as previously configured and the states of the RPV servers remain the same. Note that you can change the states of the RPV servers individually, per each server. • If you select Yes, the system automatically brings online the RPV server on the node at system restart. This puts the RPV server into the <i>available</i> state at system restart. • If you select No, the system does not bring online the RPV server at system restart. This leaves the server in the <i>defined</i> state. |

4. Press Enter. The system makes changes to the RPV servers and displays the names of the changed RPV servers and their current state: *defined* or *available*.

Removing remote physical volume servers

You can deactivate or complete delete a remote physical volume server.

To remove one or more RPV servers:

1. Enter `smit rpvserver`
2. In SMIT, select **Remove Remote Physical Volume Servers** and press Enter. A list of previously configured servers displays. Press F7 to select multiple RPV servers that you wish to remove and press Enter.
3. Enter field values as follows:

| <i>Remove Remote Physical Volume Servers fields</i> | |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Remote Physical Volume Server | The names of the RPV servers that you want to remove. |
| Keep Definitions in Database? | <p>Yes is the default.</p> <ul style="list-style-type: none"> • If you select Yes, the RPV server is deactivated but remains configured in the system, that is, remains in the <i>defined</i> state. • If you select No, the RPV server is completely deleted and its configuration is removed from the system. |

- Press Enter. SMIT displays `Are you sure?` Press Enter again. The selected RPV servers are removed.

Activating defined remote physical volume servers

You can activate one or more previously defined RPV server.

To activate one or more servers:

- Enter `smit rpvserver`
- In SMIT, select **Configure Defined Remote Physical Volume Servers** and press Enter. A list of previously configured servers displays.
- Press F7 to select multiple RPV servers that you wish to activate (or put into the *available state*) and press Enter. The system activates the RPV servers and displays their names and states, as in the following example:

```
rpvserver0 Available
rpvserver1 Available
rpvserver2 Available
```

Additional configuration tasks for RPV clients

Using the `smit rpvclient` fastpath, or the SMIT path **System Management > Devices > Remote Physical Volume Clients**, you can perform additional tasks.

Listing all remote physical volume clients

You can use SMIT to list of the remote physical volume clients.

To list all RPV clients:

- Enter `smit rpvclient`
- In SMIT, select **List All Remote Physical Volume Clients** and press Enter.

Result

The system lists all previously configured RPV clients, as in the following example:

| RPV Client | Physical Volume Identifier | Remote Site |
|------------|----------------------------|-------------|
| hdisk30 | 00010006859ece05 | central |
| hdisk32 | 0001000181810cff | central |
| hdisk51 | 0001000685a10b38 | central |

Changing or showing a remote physical volume client

You can use SMIT to change or show a remote physical volume client.

To change or show an RPV client:

- Enter `smit rpvclient`
- In SMIT, select **Change/Show a Remote Physical Volume Client** and press Enter. The system displays a list of physical disks on the remote note for which RPV clients are configured. Select an `hdisk` for which you wish to change the RPV client and press Enter.
- Enter field values as follows:

| RPV client fields | |
|-------------------------------|---------------------------------------------------------------------|
| Field | Value |
| Remote Physical Volume Client | The name of the RPV client for which you want to change parameters. |

| Field | Value |
|------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Physical Volume Identifier | The Physical Volume Identifier (PVID) associated with the RPV client. This value cannot be changed. |
| Remote Physical Volume Server Internet Address | The IP address of the RPV server to which a connection from the RPV client is allowed. Type in from one to four comma-separated new addresses in this field, or press F4 to select from the list of known IP addresses for the RPV servers defined on the remote node. |
| Remote Physical Volume Local Internet Address | The local IP addresses from which the connection to the RPV server is made. The RPV server accepts only connections from its configured RPV client IP addresses. Type in a from one to four new comma-separated addresses in this field or select an address from the picklist of known local IP addresses of existing RPV clients. |
| I/O Timeout Interval (Seconds) | How many seconds the RPV client should wait for the RPV server to respond to an I/O request before giving up and assuming that the RPV server is down on the node or is unreachable. The default is 180 seconds. In GLVM for PowerHA® SystemMirror® Enterprise Edition, this time automatically defaults to the Total Event Duration Time setting that you can specify for the config_too_long warning timer in the Change/Show Time Until Warning panel in SMIT. |

4. Press Enter. The system makes changes to the RPV client and displays the RPV client name and its current state: defined or available.

Changing multiple remote physical volume clients

You can use SMIT to change multiple remote physical volume clients.

To change or show multiple RPV clients simultaneously:

1. Enter `smit rpvclient`
2. In SMIT, select **Change Multiple Remote Physical Volume Clients** and press Enter. The system displays a list of RPV clients. Press F7 to select one or more RPV clients you want to change and press Enter.
3. Change values in some of these fields, as follows and press Enter. Any fields that are not to be changed should be blank:

| <i>Change Multiple Remote Physical Volume Clients fields</i> | |
|--------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Remote Physical Volume Clients | The names of the RPV clients for which you want to change parameters. |
| New Remote Physical Volume Server Internet Address | The service IP address of the RPV server to which a connection from the RPV client is allowed. Type in from one to four comma-separated new address in this field, or press F4 to select from the list of known IP addresses for the RPV server defined on the remote node. |
| New Remote Physical Volume Local Internet Address | The local IP address from which the connection to the server is made. The server accepts only connections from its configured client IP addresses. Type in from one to four comma-separated new address in this field or select an address from the picklist of known local IP addresses of existing RPV clients. |
| New I/O Timeout Interval (Seconds) | How many seconds the RPV client should wait for the RPV server to respond to an I/O request before giving up and assuming that the RPV server is down on the node or is unreachable. |

| Field | Value |
|-------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | The default is 180 seconds. Since you have GLVM for PowerHA® SystemMirror® Enterprise Edition installed by the time you are changing your RPV client, this time automatically defaults to the Event Duration Time setting specified for the config_too_long warning timer in the Change/Show Time Until Warning panel in SMIT. |

Removing remote physical volume clients

You can use SMIT to remove remote physical volume clients.

To remove one or more RPV clients:

1. Enter `smit rpvclient`
2. In SMIT, select **Remove Remote Physical Volume Clients** and press Enter. The system displays a list of RPV clients. Press F7 to select one or more RPV clients you want to remove and press Enter.
3. Enter field values as follows:

| <i>Remove Remote Physical Volume Clients fields</i> | |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Remote Physical Volume Clients | The names of the RPV clients which you want to remove. |
| Keep definitions in database? | <p>Yes is the default.</p> <ul style="list-style-type: none"> • If you select Yes, the RPV client is deactivated but remains configured in the system, that is, remains in the defined state. • If you select No, the RPV client is completely deleted and its configuration is removed from the system. |

4. Press Enter. SMIT displays `Are you sure?` Press Enter again. The selected RPV clients are removed.

Activating defined remote physical volume clients

Use SMIT to activate defined remote physical volume clients.

To activate one or more previously defined RPV clients:

1. Enter `smit rpvclient`
2. In SMIT, select **Configure Defined Remote Physical Volume Clients** and press Enter. A list of previously configured RPV clients displays. Press F7 to select multiple RPV clients that you want to activate (or put into the *available* state) and press Enter.

The system activates the RPV clients and displays their names and states, as in the following example:

```
hdisk1 Available
hdisk2 Available
hdisk3 Available
```

Changing RPVs and logical volume

These topics describe the maintenance tasks for geographically mirrored volume groups.

Removing remote physical volumes from a volume group

Before removing the remote physical volume from the volume group, delete all logical volume data on the physical volume. If a logical volume spans multiple physical volumes, the removal of any of the physical volumes may jeopardize the integrity of the entire logical volume.

To remove RPVs from a volume group:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Volume Groups > Remove Remote Physical Volumes from a Volume Group** and press Enter.
A list of all non-root volume groups which contain RPVs appears.
3. Select a volume group from which you want to remove an RPV and press Enter. A list of remote physical volumes in the volume group appears, as in the following example:

```
# All Remote Physical Volumes in datavg1
# RPV      PVID              Site
hdisk6    0001233459988776      site2
hdisk3    0009876352233445      site2
hdisk2    0006786544455667      site2
```

4. Select one or more RPVs you want to remove and press Enter.
5. Enter field values as follows:

| <i>Remove Remote Physical Volumes from a Volume Group fields</i> | |
|------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Volume Group Name | This is the name of the geographically mirrored volume group from which you want to remove RPVs. |
| Force | The default is No . If you select Yes , the system forces the remote physical volume to be removed from the specified volume group. |
| Remote Physical Volume Name | The name of the physical volume, such as <code>hdisk3</code> that you want to remove. More than one RPV can be removed. |

6. SMIT asks “Are you sure?”
7. Press Enter again. The GLVM uses the **reducevg** command and removes the remote physical volumes from the volume group. You can now use the **List Geographically Mirrored Volume Groups** SMIT field to check whether this volume group has any other RPVs.

Replacing a remote physical volume with another on the same site

In some cases, such as in the case of the hardware failure, it may be necessary to replace a remote physical volume with another remote physical volume on the same site.

To replace an RPV in a volume group with another RPV on the same remote site:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Volume Groups > Replace a Remote Physical Volume with Another on the Same Site** and press Enter. A list of volume groups on the remote site is shown.
3. Select a volume group from the picklist and press Enter. A list of RPVs configured on the remote site is shown.
4. Select one or more RPVs to remove and press Enter. A list of RPVs is shown.

- Select one or more RPVs to add and press Enter. The **Replace a Remote Physical Volume** screen is shown.
- Enter field values as follows:

| <i>Replace a Remote Physical Volume fields</i> | |
|------------------------------------------------|------------------------------------------------------------------------------------------------|
| Field | Value |
| Force | The default is No . |
| Old Remote Physical Volume Name | The name of the physical volume, such as <code>hdisk3</code> that you want to replace. |
| New Remote Physical Volume Name | The name of the remote physical volume with which you want to replace the old physical volume. |

- Press Enter. The GLVM runs the `replacepv` command to replace the RPV.

Listing geographically mirrored volume groups

You can use SMIT to list geographically mirrored volume groups.

To list all geographically mirrored volume groups that are already configured:

- Enter `smit`
- In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Volume Groups > List Geographic Volume Group Information** and press Enter.

Result

SMIT displays geographically mirrored volume groups that are configured, as in the following example:

| Volume Group | Logical Volume | RPV | PVID | Site |
|--------------|----------------|--------|------------------|---------|
| vg1 | lv1 | hdisk3 | 0001000181810988 | Central |

Removing a remote site mirror copy from a logical volume

You can deallocate mirror copies from each logical partition in the logical volume.

To deallocate mirror copies:

- Enter `smit`
- In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Remove a Remote Site Mirror Copy from a Logical Volume** and press Enter.
SMIT displays all logical volumes that have mirrors created at a remote site.
- Select a logical volume for which you want to remove a mirror copy at a remote site and press Enter. If more than one mirror copy exists, SMIT displays a list of mirror copies. If only one mirror copy exists, the GLVM utilities do not allow you to remove it and issue an error.
- Select a mirror copy which you want to remove and press Enter. SMIT displays the details of the logical volume for which you want to remove a mirror copy.
- Enter or review field values as follows:

| <i>Remove a Remote Site Mirror Copy from a Logical Volume fields</i> | |
|----------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Logical Volume Name | The name of the selected logical volume. |
| New Total Number of Logical Partitions Copies | Specifies the maximum number of logical partition copies. Copies are the physical partitions which, in addition to the original physical partition, make up a logical partition. You can have up to three copies in a logical volume. |

| Field | Value |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Remote Physical Volume Names | Select the remote physical volume name. This is the name of the remote physical volume that belongs to the logical volume for which you want to remove a copy. Remote physical volume names are typically in the form "hdiskx" where x is a unique number system-wide. |

6. Press Enter. The SMIT asks "**Are you sure?**" GLVM removes one copy at a time. If a mirror copy of the logical volume spans more than one disk, the entire mirror copy is removed. The system uses the RPV to identify which copies to remove, if there is more than one copy.

Listing all geographically mirrored logical volumes

You can use SMIT to list all geographically mirrored logical volumes that are already configured.

To list all volumes:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > List Geographic Logical Volume Information** and press Enter.

Result

SMIT displays all remote physical volumes and their assignment to volume groups and logical volumes in the following format:

| # | Volume Group | Logical Volume | RPV | PVID | Site |
|-----|--------------|----------------|--------------|------|---------|
| vg1 | lv1 | hdisk1 | 000987654321 | | chicago |

Verifying mirror copy site locations for a logical volume

With this SMIT option, you verify that logical volume mirrors which have geographic copies have been configured correctly by determining that each copy has physical volumes that reside on only one site and that the logical volume has the Super Strict inter-disk allocation policy.

To verify configuration:

1. Enter `smit`
2. In SMIT, select **System Storage® Management (Physical and Logical Storage) > Geographic Logical Volume Manager Utilities > Geographically Mirrored Logical Volumes > Verify Geographic Logical Volume Information** and press Enter.
SMIT displays a list of logical volumes.
3. Select the logical volumes from the list, or enter the names of the logical volumes for which you want to verify geographic mirrors and press Enter.
When you verify the mirror copies, the system verifies that:
 - A copy of the logical volume exists on only one site: A local copy exists totally on the local site, and the remote copy totally exists on the geographically remote site.
 - All logical volumes are configured with the Super Strict inter-disk allocation policy.

Result

When you select the logical volumes from the list, or enter the names of the logical volumes and press Enter, the system issues a message: "Checking logical volume <lv_name>". After a brief delay, the SMIT interface displays either a successful command completion or issues an error.

Converting a standard cluster to a linked cluster

In PowerHA® SystemMirror® Version 7.2.6, or later, existing users can migrate a data center from the local site to the cloud environment to use the Geographic Logical Volume Manager (GLVM) features.

To convert an existing standard cluster to a linked cluster, add new node at the remote site (secondary site) and add all the existing cluster nodes at the local site (primary site). You must stop all cluster services during cluster conversion.

The following figure shows an example conversion of a standard cluster to a linked cluster. The standard cluster in the following example consists of two nodes, NodeA, NodeB, and repository disk hdisk1. After conversion, a linked cluster is created which consist of an existing standard cluster at the local site, and a new node NodeC and a new repository disk hdisk2 at the remote site.

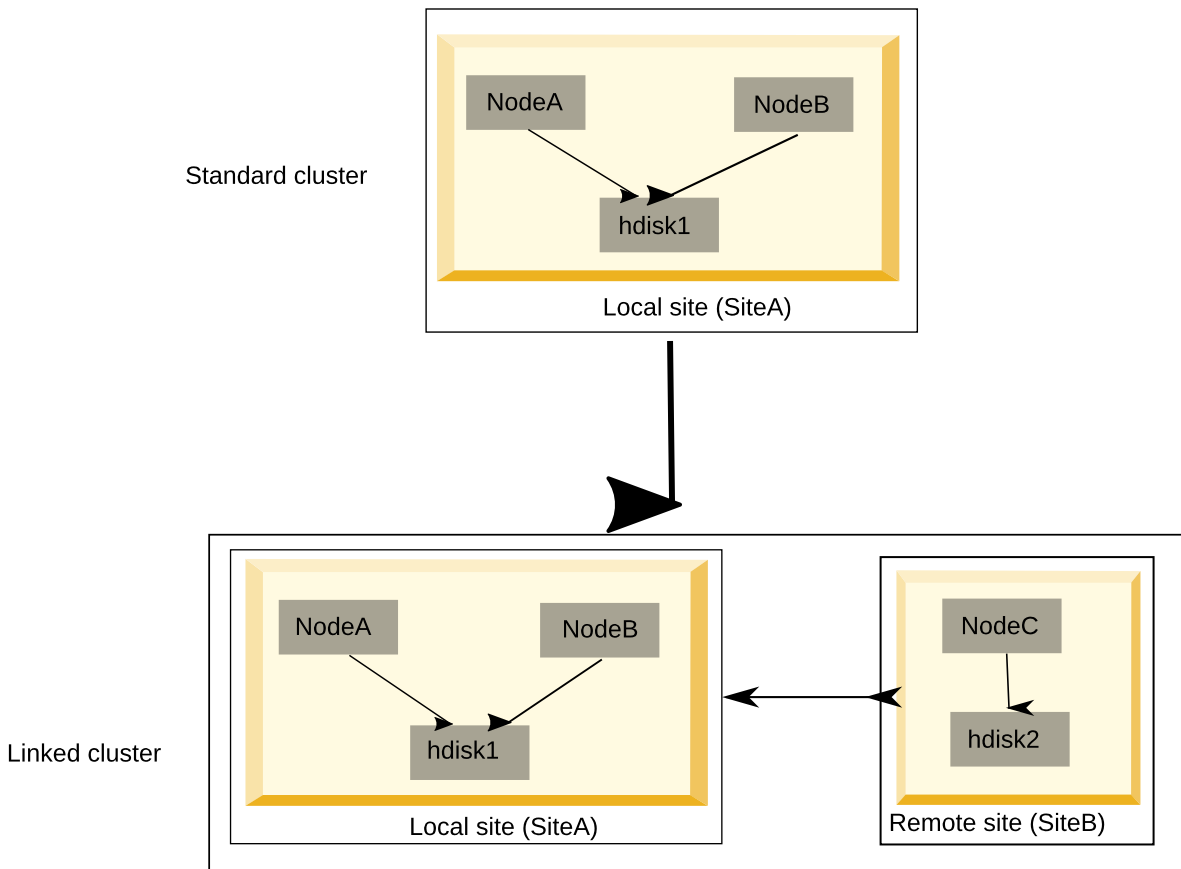


Figure 1: Standard cluster to a linked cluster conversion

To convert an existing cluster to a linked cluster, complete the following steps:

1. Take a snapshot of an existing standard cluster and save it. For instruction about creating a snapshot of the cluster, see [Creating a snapshot of the cluster configuration](#).
2. To change the standard cluster type, run the following command.

```
# clmgr modify cluster TYPE=LC
```

3. To create a local site and to add all nodes of the standard cluster to it, run the following command:

```
# clmgr add site <local site name> NODES=<All nodes part of existing standard cluster>
```

For example, to create a local site, SiteA and to add two cluster nodes, NodeA and NodeB, run the following command:

```
clmgr add site SiteA NODES=NodeA,NodeB
```

4. To create a remote site with a new node and a new repository disk, run the following command:

```
# clmgr add site <remote site name> NODES=<Remote site node>  
REPOSITORIES=<remote site repository disk name>
```

For example, to create a remote site, SiteB with a new node NodeC, and a new repository disk hdisk2, run the following command:

```
clmgr add site SiteB NODES=NodeC REPOSITORIES=hdisk2
```

5. To synchronize the linked cluster, run the following command:

```
# clmgr sync cluster
```

- To verify conversion of the standard cluster to a linked cluster, run the following command:

```
# clmgr query cluster
```

Note: In case of conversion failure, restore the previous cluster configuration. For instruction, see [Restoring the cluster configuration from a snapshot](#).

Integrating geographically mirrored volume groups into a PowerHA® SystemMirror® cluster

These topics describe how to configure a resource group in the PowerHA® SystemMirror® cluster that includes two sites and geographically mirrored volume groups.

Identifying the configuration path

The configuration path you select depends on the current hardware and software installation you have. This section lists the possible scenarios and provides instructions for each. After identifying the applicable scenario, read only the sections describing it.

Depending on the software and hardware you currently have installed, you can have one of the following cases:

| Your Current® Installation | Your goal with using GLVM for PowerHA® SystemMirror® Enterprise Edition | What You Do Next |
|---------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PowerHA® SystemMirror® cluster configured on a single site | Start using GLVM for PowerHA® SystemMirror® Enterprise Edition as a disaster recovery solution for applications that are currently supported by PowerHA® SystemMirror® on a single site. | See “Installation options” on page 36 . After you install GLVM for PowerHA® SystemMirror® Enterprise Edition, start with “Geographically mirrored volume groups configuration prerequisites” on page 80 . |
| A set of nodes and PowerHA® SystemMirror® clusters not defined yet. | Start using GLVM for PowerHA® SystemMirror® Enterprise Edition as your disaster recovery solution for applications that you want to keep highly available. | After you install and plan GLVM for PowerHA® SystemMirror® Enterprise Edition, start with “Geographically mirrored volume groups configuration prerequisites” on page 80 . |

Geographically mirrored volume groups configuration prerequisites

Before configuring, ensure that you have completed all the tasks in the previous sections:

- Installed GLVM for PowerHA® SystemMirror® Enterprise Edition on all nodes at each site.
- Planned and configured the physical volumes, volume groups and logical volumes on the nodes at the local site using the standard AIX® SMIT interface.

Note: You can configure enhanced concurrent mode volume groups (that are also geographically mirrored) only if the PowerHA® SystemMirror® cluster is configured and the PowerHA® SystemMirror® cluster services are running on the nodes. Therefore, prior to configuring enhanced concurrent volume groups with RPVs, ensure that you start the cluster services on the nodes at the local (primary) site.

- Planned and configured sites. If you have PowerHA® SystemMirror® installed, but have not used sites, add one local and one remote site to the cluster, ensuring that at least one node at each geographic location belongs to the site and that the inter-site management policy is not set to Ignore or to Online on both Sites.

- Configured RPV server and client pairs on all nodes at each site for all volume groups that you plan to include in resource groups. Site names in PowerHA® SystemMirror® must match the site names defined for RPV servers. After you complete the configuration of RPV servers and clients for volume groups, RPV servers and clients must be in the defined state.
- Added mirror copies to the volume groups and logical volumes using the **GLVM Utilities** SMIT interface.
- Planned for PowerHA® SystemMirror® Enterprise Edition networks.

Management of RPVs in a PowerHA® SystemMirror® cluster

After RPVs become part of volume groups that belong to an active PowerHA® SystemMirror® cluster, PowerHA® SystemMirror® takes care of switching the RPV states between *defined* and *available* during cluster events.

If, for some reason, you must stop the cluster services and manage the RPVs manually, the RPVs must be configured in the *defined* state before you start the cluster services. This enables management of RPVs by PowerHA® SystemMirror® after the nodes join the cluster.

Overview of configuration tasks

To set up a GLVM for PowerHA® SystemMirror® Enterprise Edition configuration, there are several steps you will need to take.

These steps include:

1. [“Configuring sites” on page 81](#)
2. [“Configuring XD-type networks” on page 81](#)
3. Define GLVM volume groups before adding them to resource groups.
Convert the existing regular volume groups that are part of your resource groups into geographically mirrored volume groups.

For enhanced concurrent volume groups that you want to also make geographically mirrored, you need to have PowerHA® SystemMirror® cluster services running, before adding and managing RPVs in such groups.

4. [“Preparing the geographically mirrored volume groups to be put under PowerHA SystemMirror control” on page 84](#)
5. [“Configuring resource groups in GLVM for PowerHA SystemMirror Enterprise Edition” on page 83](#)

Configuring sites

To properly configure sites to use in GLVM for PowerHA® SystemMirror® Enterprise Edition, there are several steps to take.

To configure sites to use in GLVM for PowerHA® SystemMirror® Enterprise Edition, complete the following steps:

1. Configure one local and one remote site.
2. Match the site names to the site names defined for RPVs.
3. Select any inter-site policy except **Ignore** and **Online on Both Sites**.
4. Include all nodes at the same geographic location in the cluster to a site at that location.

Configuring XD-type networks

In GLVM for PowerHA® SystemMirror® Enterprise Edition, you can include up to four XD_data networks in the PowerHA® SystemMirror® cluster.

In previous releases of GLVM, no IP address takeover is supported; service IP labels are node-bound. You can continue to use this configuration. With GLVM for PowerHA® SystemMirror® Enterprise Edition, XD_data networks have IPAT via IP aliases enabled by default. With IPAT via IP aliases enabled, you configure persistent node IP labels.

XD_data networks with IPAT via IP aliases enabled can also have a multiple-node alias service IP label for clients to access. You can make this a site-specific service IP label.

To prevent cluster partitioning, add an XD_ip network.

Note: NFS cross-mounts are not supported on XD networks.

Configuring GLVM to use XD networks after installing PowerHA® SystemMirror®

After you install PowerHA® SystemMirror®, you still have to configure GLMV to use XD networks.

To configure GLMV to use XD networks, complete the following steps.

1. Using PowerHA® SystemMirror® SMIT, create the XD_data network.
2. For each node, create one persistent IP label and add it to the XD_data network.
3. Verify and synchronize the PowerHA® SystemMirror® cluster. This will activate the persistent IP labels.
4. Configure GLVM to communicate over the persistent IP labels.

Configuring an XD_data network

Use the SMIT interface to configure an XD_data network.

To configure an XD_data network, complete the following steps:

1. Enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Cluster Nodes and Networks > Initial Cluster Setup (Custom) > Networks->Change/Show a Network**, and press Enter.
3. SMIT displays a list of networks you can configure. Select the network you want to change and on the next panel change the **Network Type** field to **XD-data**.
4. Press Enter to configure the XD_data network.

Configuring persistent IP labels on XD_data networks

You can use SMIT to configure persistent IP labels on XD_data networks.

To configure a persistent IP label on the XD_data network:

1. Add an IP-based network to the PowerHA® SystemMirror® cluster that has IP Takeover via IP aliases enabled.
2. Add a persistent IP Label/Address to the XD_data network for each node by entering `smit sysmirror` from the command line.
3. From the SMIT interface, select **Cluster Nodes and Networks > Manage Nodes > Configure Persistent Node IP Label/Addresses > Add a Persistent Node IP Label/Address**
4. Select the node where you want to add the persistent IP label/address.
5. Enter the field values as follows:

| <i>Add Persistent Node IP Label/Address fields</i> | |
|----------------------------------------------------|----------------------------------------------------------------------|
| Field | Value |
| Node Name | The name of the node on which the IP label/address will be bound. |
| Network Name | The name of the network on which the IP label/address will be bound. |
| Node IP Label/Address | The IP label/address to keep bound to the specified node. |

6. Press Enter.
7. Repeat the steps to add a persistent IP label/address to each node.

Configuring XD_ip networks

An XD_ip network is an IP-based network that is used for client communication. The network used for the XD_ip network type is either too slow or does not have enough bandwidth to be suitable for data mirroring. However, you can configure this network for heartbeating and to prevent cluster partitioning.

Heartbeating over disks which are members of GLVM volume groups is supported, but only within each site using disks local to that site which are part of enhanced concurrent mode volume groups. Heartbeating between the sites is done through the XD_ip network.

This network uses IPAT via IP aliases by default. You must add at least one service IP label to this network, and use regular service IP labels.

To configure an XD_ip network, complete the following steps:

1. Enter `smit sysmirror`
2. In SMIT, select **Custom Cluster Configuration > Cluster Nodes and Networks > Initial Cluster Setup (Custom) > Networks->Change/Show a Network**, and press Enter. SMIT displays a list of network types.
3. From the list of networks, select the one you wish to change, and change the **Network Type** field to **XD_ip**.
4. Press Enter to configure this network.
5. Create IP labels on the network.

Configuring site-specific IP labels

Service IP labels can have a site association. Site-specific service IP labels are activated only when their resource group is set to **Online Primary** on the specified site.

The service IP label can failover to other nodes on the specified site. These service IP labels function the same as regular service IP labels, except for the ability to use NFS cross-mounts and the site-specific information.

On an XD_data network, configure a service IP label, and add it to a resource group that includes nodes for the required site.

Verifying XD_data network configurations

The verification utility checks the XD_data network configuration.

The verification utility checks the XD_data network for the following configurations:

- A persistent node IP label exists for each node on the network.
- A service IP label is not configured on multiple nodes because the increase in network traffic might interfere with the GLVM data replication traffic. If possible, use an XD_ip network instead.

Configuring resource groups in GLVM for PowerHA® SystemMirror® Enterprise Edition

After defining RPVs and geographically mirrored volume groups and logical volumes using the **GLVM Utilities** SMIT interface, you can add the volume groups to the resource groups.

Prerequisites and limitations

The following restrictions apply:

- You cannot include a geographically mirrored volume group into a resource group when the RPV server is hosted on a node outside of the PowerHA® SystemMirror® cluster.
- You cannot use System Management (C-SPOC) panels in the SMIT interface for creation and management of geographically-mirrored volume groups.
- In order to configure enhanced concurrent volume groups with RPVs (and thus make them geographically mirrored) you must have PowerHA® SystemMirror® cluster services running on the nodes on which you perform the operations.

- The field **Use Forced Varyon of Volume Groups, if Necessary** should be set to **true**, although the default is **false**. This setting allows GLVM for PowerHA® SystemMirror® Enterprise Edition to vary on the volume group if quorum is disabled and the remote site fails.
- The field **Automatically Import Volume Groups** should be set to **false**. Even if you set this field to **true**, PowerHA® SystemMirror® does not automatically import geographically mirrored volume groups and issues a warning. If **Automatically Import Volume Groups** is set to **false**, PowerHA® SystemMirror® does not import the volume groups automatically. In this case, when you add volume groups to the resource group, make sure that the volume groups have already been imported to each of the nodes using the **importvg** command (C-SPOC is not supported in GLVM for PowerHA® SystemMirror® Enterprise Edition).
- You can configure nonconcurrent and concurrent resource groups that will contain geographically mirrored volume groups. The policies are already filled in, in the SMIT screens.
- For nonconcurrent resource groups, the startup policy could be any of the nonconcurrent startup policies, such as Online on First Available Node, or Online on Home Node Only; the fallover policy is Fallover to Next Priority Node in the List and Fallback could be either Fallback to Higher Priority Node in the List or Never Fallback.
- For concurrent resource groups, the startup policy is Online on All Available Nodes, the fallover policy is Bring Offline (on Error) and the fallback policy is Never Fallback.

Preparing the geographically mirrored volume groups to be put under PowerHA® SystemMirror® control

After configuring geographically mirrored volume groups, perform some steps to prepare the volume groups for PowerHA® SystemMirror® cluster verification.

The goal of these steps is to ensure that all geographically mirrored volume groups are imported on all cluster nodes that belong to the resource group's node list.

- These steps are relevant for standard and enhanced concurrent mode volume groups.
- These steps ensure that before you transfer control over the geographically mirrored volume group to PowerHA® SystemMirror®, the volume group is varied off, and its associated RPV servers and clients are configured but inactive (remain in the *defined* state). After PowerHA® SystemMirror® establishes control over the geographically mirrored volume group, it activates the volume group and its RPV servers and clients as necessary, during the cluster events.
- It is important to perform these steps prior to adding the volume group to a resource group. The PowerHA® SystemMirror® cluster verification enables PowerHA® SystemMirror® to automatically include the geographically mirrored volume group into the picklist of volume groups that you can add to a resource group.

To prepare a geographically mirrored volume group for PowerHA® SystemMirror® cluster verification:

1. On the local node, vary off the volume group.
2. On the local node, change the RPV clients to the *defined* state.
3. On the remote node, change the RPV servers to the *defined* state.
4. On the local node, create RPV servers for each local physical volume (PV) that belongs to the geographically mirrored volume group, and leave the RPV servers in the *available* state.
5. On the remote node, create RPV clients for each local PV that belongs to the geographically mirrored volume group, and leave the RPV clients in the *available* state.
6. Import the geographically mirrored volume group on the remote node.
7. In the AIX® SMIT path, **System Management > System Storage® Management (Physical and Logical Storage) > Logical Volume Manager > Volume Groups > Set Characteristics of a Volume Group > Change a Volume Group**, select the volume group and set the **Activate a Volume Group Automatically** field to **No**.
8. In the SMIT fastpath, `smit rpvserver` select **Change/Show a Remote Physical Volume Server**. Select the RPV server and set the field **Configure Automatically at System Restart?** to **No**. This ensures that the system *does not* automatically bring online the RPV server on the node at system restart, and that the RPV server remains in the *defined* state at system restart.
9. Vary off the volume group on the remote node.

10. On the remote node, change the RPV clients to the *defined* state.
11. On the local node, change the RPV servers to the *defined* state.
12. Repeat steps 4 - 11 as necessary for all nodes in the associated resource group (that is, the group into which you plan to include the volume group).

Result

The volume group is now ready to be included in a PowerHA® SystemMirror® resource group, and should pass cluster verification. To continue, use the procedure in Steps to add geographically mirrored volume groups to resource groups.

Sample procedure for Preparing a volume group for inclusion in the PowerHA® SystemMirror® resource group

This sample procedure provides you will information about preparing a volume group for inclusion in the PowerHA® SystemMirror® cluster group.

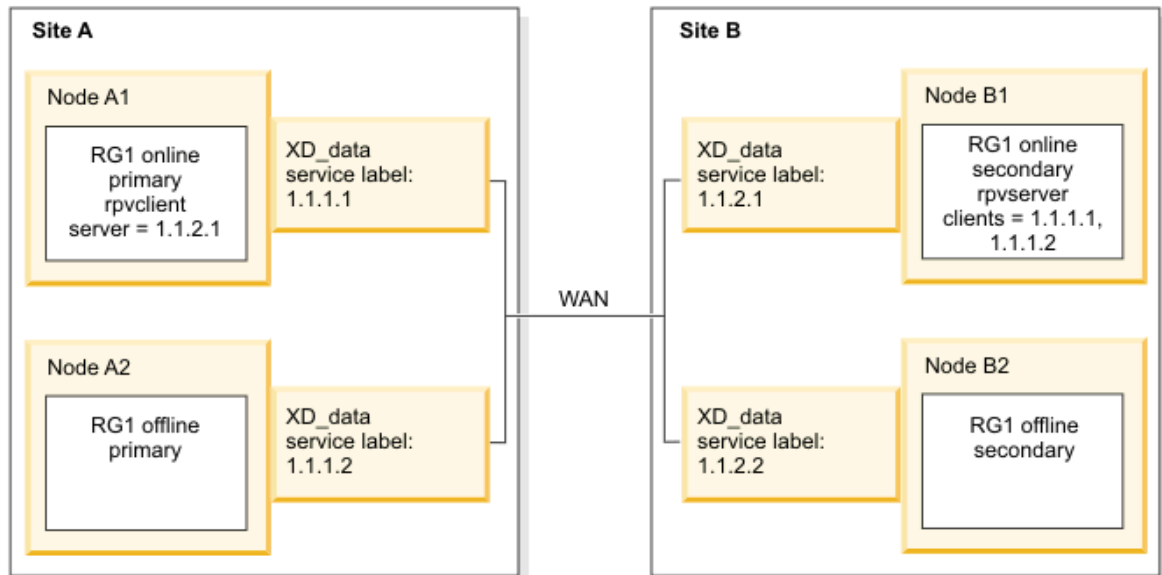
Consider the following sample cluster configuration.

- Local site A with nodes A1 and A2.
- Remote site B with nodes B1 and B2.

The following figure illustrates the configuration.

Note: For the resource groups to be in the states in which they are shown in the figure, the tasks in the rest of this section must be completed. Also, if your geographically mirrored volume groups also happen to be *enhanced concurrent mode* volume groups, then RG1 would be in the ONLINE PRIMARY state on node A2, in addition to being ONLINE PRIMARY on node A1.

If you have more than one XD_data network, then multiple RPV servers would be added on different nodes at site B, not just one, as shown. To simplify the diagram, however, multiple networks and RPVs are not displayed.



To fully prepare a previously configured geographically mirrored volume group for inclusion in the PowerHA® SystemMirror® resource group.

| Task | Task Details |
|-------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create a geographically mirrored volume group on node A1 (site A) | On nodes B1 and B2, configure the RPV servers, and change the server on node B1 to the <i>available</i> state. On node A1, configure the RPV client, and change it to the <i>available</i> state. |

| Task | Task Details |
|----------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | On node A1, configure the volume group VG1, mirrored logical volumes and file systems. |
| Import the geographically mirrored volume group on node A2 (site A) | <p>On node A1, unmount file systems, vary off the volume group and change the RPV client to the <i>defined</i> state.</p> <div style="border: 1px solid blue; padding: 5px; margin: 10px 0;"> <p>Note: For enhanced concurrent volume groups, on node A1, do <i>NOT</i> unmount file systems, vary off the volume group and change the RPV client to the <i>defined</i> state.</p> </div> <p>On node A2, configure the RPV client, and change it to the <i>available</i> state.</p> <p>Import the volume group on node A2.</p> |
| Import the volume group on node B1 (site B) | <p>On node A2, unmount file systems, vary off the volume group, change the RPV client to the <i>defined</i> state.</p> <div style="border: 1px solid blue; padding: 5px; margin: 10px 0;"> <p>Note: For enhanced concurrent volume groups, vary off the volume group on <i>all</i> nodes at site A.</p> </div> <p>On node B2, change the RPV server to the <i>defined</i> state.</p> <p>On nodes A1 and A2, configure the RPV servers, and move the server on node A1 to the <i>available</i> state.</p> <p>On node B1, configure the RPV client, and move it to the <i>available</i> state.</p> <p>On node B1, import the volume group.</p> |
| Import the geographically mirrored volume group on node B2 (site B) | <p>On node B1, unmount file systems, vary off the volume group, and change the RPV client to the <i>defined</i> state.</p> <p>On node B2, configure the RPV client, and move it to the <i>available</i> state.</p> <p>On node B2, import the volume group.</p> <p>On node B2, unmount file systems, vary off the volume group, and change the RPV client to the <i>defined</i> state.</p> |

Adding geographically mirrored volume groups to resource groups

Use SMIT to add geographically mirrored volume groups to the resource groups.

To add geographically mirrored volume groups to the resource groups:

1. Enter `smit sysmirror`.
2. In SMIT, select **Custom Cluster Configuration > Resource Groups->Change/Show Resources and Attributes for a Resource Group**, and press Enter.
SMIT displays a list of defined resource groups.
3. Select the resource group you want to configure and press Enter. SMIT returns the panel that matches the type of resource group you selected, with the **Resource Group Name**, **Inter-site Management Policy**, and **Participating Node Names (Default Node Priority)** fields filled in.
SMIT displays only valid choices for resources, depending on the resource group startup, failover and fallback policies that you selected. For a list of allowed policies see Prerequisites and limitations.

Be aware that once you add resources to a resource group, its startup, fallover and fallback policies cannot be changed unless you remove the resources. You can only change the resource group policies in a resource group that does not contain any resources yet. Plan your resource group's policies in advance, before adding resources to it.

If the cluster services are running on the participating nodes, press F4 to list the shared resources. If a resource group and node relationship has not been defined, or if the cluster services are not running on a node, F4 displays the appropriate warnings.

In the following step, not all SMIT fields display. This procedure only contains the SMIT fields that are relevant to the inclusion of geographically mirrored volume groups into the resource group.

4. Enter the field values as follows. A nonconcurrent resource group is shown; however, concurrent resource groups can also be configured.

| <i>Resource group fields</i> | |
|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Dynamic Node Priority (Overrides default) | Select the dynamic node priority policy. The default is blank (the ordered node list). The preconfigured dynamic node priority policies are listed. |
| Service IP Labels/ Addresses | <p>List the IP labels/addresses to be taken over when this resource group is acquired. Press F4 to see a valid list. The list includes addresses which rotate or may be taken over.</p> <div data-bbox="564 1043 1426 1285" style="border: 1px solid black; padding: 5px;"> <p>This field is not relevant for the task of adding geographically mirrored volume groups: You cannot include service IP labels from the XD_data network in the resource group, since these labels are node-bound and remain on the node, once configured. This field is intended for service IP labels that are associated with the application that you want to keep highly available.</p> </div> |
| Application Controllers | Indicate the application controllers to include in the resource group. Press F4 to see a list of application controllers. |
| Volume Groups | <p>Identify the shared volume groups that should be varied on when this resource group is acquired or taken over. Select the volume groups from the picklist or enter desired volume groups names in this field.</p> <p>If you previously requested that PowerHA® SystemMirror® collect information about the appropriate volume groups, then pressing F4 will give you a list of all shared volume groups in the resource group and the volume groups that are currently available for import onto the resource group nodes.</p> <p>Specify the shared volume groups in this field if you want to leave the field Filesystems (empty is All for specified VGs) blank and to mount all file systems in the volume group. If you specify more than one volume group in this field, then all file systems in all specified volume groups will be mounted; you cannot choose to mount all file systems in one volume group and not to mount them in another.</p> <p>For example, in a resource group with two volume groups, vg1 and vg2, if the Filesystems (empty is All for specified VGs) is blank, then all the file systems in vg1 and vg2 will be mounted when the resource group is brought up. However, if the Filesystems (empty is All for specified VGs) has only file systems that are part of the vg1 volume group, then none of the file systems in vg2 will be mounted, because they were not entered in the Filesystems (empty is All for specified VGs) field along with the file systems from vg1.</p> |

| Field | Value |
|--------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | If you have previously entered values in the Filesystems field, the appropriate volume groups are already known to the PowerHA® SystemMirror® software. |
| Use Forced Varyon of Volume Groups, if Necessary | For geographically mirrored volume groups, set this flag to true . (The default is false). If this flag is set to true , PowerHA® SystemMirror® uses a forced varyon to bring each volume group that belongs to this resource group online in the event that a normal varyon for the volume group fails due to a lack of quorum, and if PowerHA® SystemMirror® finds at least one complete copy of every logical partition in every logical volume available for this volume group. This option should only be used for volume groups in which every logical volume is mirrored. It is recommended to use the super strict disk allocation policy; forced varyon is unlikely to be successful for other choices of logical volume configuration. |
| Raw Disk PVIDs | In GLVM for PowerHA® SystemMirror® Enterprise Edition, you cannot use RPVs as raw disk PVIDs. |
| Automatically Import Volume Groups | Specifies whether PowerHA® SystemMirror® should <i>automatically</i> import those volume groups that are defined in the Volume Groups or Concurrent Volume Groups fields. By default, Automatically Import Volume Groups flag is set to false . <div style="border: 1px solid blue; padding: 5px; margin: 5px 0;"> <p>Even if you set this field to true, PowerHA® SystemMirror® does not automatically import geographically mirrored volume groups and issues a warning.</p> </div> <p>If Automatically Import Volume Groups is set to false, selected volume groups are not imported automatically. In this case, when you add volume groups to the resource group, make sure that the selected volume groups have already been imported to each of the nodes using the importvg command (C-SPOC is not supported in GLVM for PowerHA® SystemMirror® Enterprise Edition).</p> |
| Allow varyon with missing data updates? (Asynchronous GLVM Mirroring Only) | Value used to determine whether PowerHA® SystemMirror® event processing should automatically allow data divergence to occur during a non-graceful site failover. True allows data divergence to occur while False prevents it. |
| Default choice for data divergence recovery (Asynchronous GLVM Mirroring Only) | Specifies your default choice for data divergence recovery ahead of time. The value of this field is the PowerHA® SystemMirror® site name of the site whose version of the data that you are most likely to choose to preserve when recovering from data divergence |

5. Press Enter to add the values to the PowerHA® SystemMirror® Configuration Database.
6. Return to the **Custom Configuration** menu and synchronize the cluster.

Converting existing volume groups into geographically mirrored volume groups

This scenario applies to you only if before planning to use GLVM for PowerHA® SystemMirror® Enterprise Edition you were using PowerHA® SystemMirror® and have volume groups that you would like to mirror remotely using the mirroring function of the GLVM.

In this case, the process for planning your geographically mirrored volume groups and configuring them as part of resource groups includes the following additional steps.

Note: In order to turn enhanced concurrent volume groups into geographically mirrored volume groups, PowerHA® SystemMirror® must be installed and the cluster services must be running, that is, you cannot configure enhanced concurrent volume groups with RPVs in them if you only have AIX® GLVM utilities installed, or if the cluster services are not running on the nodes. For instructions, see *Configuring RPV client-server pairs for an enhanced concurrent volume group on a node at each site and Extending an existing enhanced concurrent volume group to nodes at remote site*.

To convert existing standard volume groups in the cluster into geographically mirrored volume groups that are part of the resource groups:

1. Stop the cluster services on all nodes. Ensure that the volume groups are varied off.
2. In AIX® SMIT, perform the necessary steps to create RPV servers and RPV clients on each node and each site for the volume groups that you plan to mirror remotely.
3. In AIX® SMIT, create the mirror copies of logical volumes and volume groups.
4. Verify the configuration using the **GLVM Utilities** SMIT interface.
5. Perform the steps in the section *Preparing the geographically mirrored volume groups to be put under PowerHA® SystemMirror® control*.
6. Start the cluster services.
7. Configure XD-type networks.
8. Configure sites in PowerHA® SystemMirror®.
9. Add the geographically mirrored volume groups as resources to the existing resource groups.

Verifying and synchronizing the GLVM configuration in PowerHA® SystemMirror®

The configuration changes that you have completed to this point need to be synchronized to the other cluster nodes.

Note: When you configure geographically mirrored volume groups, some automatic corrective actions of the PowerHA® SystemMirror® cluster verification utility, such as automatically importing or exporting a volume group, do not work. Therefore, run the automatic corrective action of the cluster verification utility in the Interactive mode.

GLVM for PowerHA® SystemMirror® Enterprise Edition verifies the configuration and ensures that the following conditions are met. Otherwise, the system issues errors. For some checks, GLVM for PowerHA® SystemMirror® Enterprise Edition issues a warning. This is indicated in parentheses.

- The filesset **cluster.xd.glvm** is installed on all cluster nodes.
- The RPV site name of the associated geographically mirrored volume group for each node matches its PowerHA® SystemMirror® site name.
- From one up to four XD_data networks are defined.
- A XD_ip network is defined to protect against site isolation. GLVM for PowerHA® SystemMirror® Enterprise Edition issues a warning if these networks are not configured.
- There is at least one node at each site in the resource group.
- All nodes in the resource group have service labels defined on the XD_data network.
- The inter-site management policy is not Ignore and is either Online on Either Site or Prefer Primary Site.
- The resource group policy could be either concurrent or nonconcurrent.
- Each node in the resource group has an RPV server defined for every local physical volume in its geographically mirrored volume group.

- If there are any geographically mirrored volume groups in the resource group, all volume groups in the resource group are geographically mirrored volume groups.
- All logical volumes in each PowerHA® SystemMirror® managed geographically mirrored volume group have the disk allocation policy set to Super Strict.
- Each PowerHA® SystemMirror® managed geographically mirrored volume group has a complete mirror copy of all logical volumes at both sites.
- PowerHA® SystemMirror® also performs all the existing checks that verify the consistency of the disks, volume groups and file systems configuration across the cluster.
- In addition, PowerHA® SystemMirror® performs the existing checks that verify the configuration of the XD_data and XD_ip networks.

Changing GLVM-mirrored resources configuration

You can change the configuration of GLVM-mirrored resources.

For information about changing the sites, IP networks and other networks in PowerHA® SystemMirror®, see *Administering PowerHA® SystemMirror®*.

For information about changing the RPV servers and clients outside of PowerHA® SystemMirror® (when the cluster services are down on the nodes), see *Configuring geographically mirrored volume groups*.

If sites are defined in the PowerHA® SystemMirror® cluster, you cannot dynamically change the cluster resources. Therefore, in GLVM for PowerHA® SystemMirror® Enterprise Edition, you cannot dynamically apply changes to the resource groups that contain geographically mirrored volume groups.

To change the geographically mirrored volume groups, stop the cluster services on all nodes, and perform the changes outside of PowerHA® SystemMirror® on each node in the cluster. After the changes are complete, use the procedure in *Preparing the geographically mirrored volume groups to be put under PowerHA® SystemMirror® control before reimporting the changed volume groups to the resource groups in the cluster*.

Adding RPV disks in a running cluster

To add a disk to a geographically mirrored volume group in a running cluster, complete the following steps. Before following the steps however make sure that all of the disks have PVIDs assigned. If PVIDs are not assigned you must use the PVIDs using the `chdev -a pv=yes -l hdisk3` (where `hdisk3` was a new disk added to the system).

Steps for adding RPV disk to a GMVG in a running cluster are divided into different sets or blocks.

Use set 1 if you want to add an RPV disk to a site holding the primary instance of RG (for example, where the volume group is varied online and applications are running).

Use set 2 If you want to add an RPV disk at the remote site (where the volume group is offline or holding the secondary instance of RG). To add a new disk at both the sites complete both procedures.

Procedure 1: Adding an RPV disk at the primary site.

1. On the node that has the volume group varied on, identify a new physical disk that will be added to the volume group. This disk must be accessible by all nodes at that site.
2. On the node that has the volume group varied on, create RPV server instance for this disk. The **Configure Automatically at System Restart** field should be set to **No**. The field **Start New Device Immediately** should be set to **Yes**.
3. Create the RPV server instance on all of the remaining nodes of this site using the disk. The field **Configure Automatically at System Restart** should be set to **No**. The field **Start New Devices Immediately** should be set to **No**.
4. On the remote site node create an RPV client for the disk. Associating it with the RPV server instance created in prior step. The field **Start New Devices Immediately** should be set to **No**.
5. Create an RPV client instance on all remaining nodes on the remote site. The field **Start New Devices Immediately** should be set to **No**.
6. On the node having volume group varied on run the `extendvg` command from the command line to include the local physical volume in the volume group.

Procedure 2: Steps for adding a disk at the backup site

1. On one of the remote site node (holding secondary instance of the Resource Group), identify a new physical disk that will be added to the volume group. This disk must be accessible by all nodes at that site.
2. Create an RPV server instance for RPV disk. The **Configure Automatically at System Restart** should be set to **No**. The **Start New Devices Immediately** field should be set to **No**.
3. Create RPV server instances on all other nodes of the remote site. The **Configure Automatically at System Restart** should be set to **No**. The field **Start New Devices Immediately** should be set to **No**.
4. On the node that has the volume group varied on (primary instance of Resource Group) create RPV client instance for the disk. The **Start New Devices Immediately** should be set to **Yes**.
5. RPV client on all other nodes on this site. The **Start New Devices Immediately** should be set to **No**.
6. On the node that has the volume group varied on, (holding the primary instance of the Resource Group) use the **extendvg** command to extend the volume group to add or include the remote physical volume or disk to the volume group.
7. Update the volume group definition on all of the other nodes at each site, enter `smit sysmirror` from the command line.
8. From the SMIT interface, select **System Management (C-SPOC) > Storage > Volume Groups > Synchronize a Volume Group Definition**, and press Enter.

Managing GLVM for PowerHA® SystemMirror® Enterprise Edition

This section describes situations you might encounter in a running cluster after synchronizing your GLVM for PowerHA® SystemMirror® Enterprise Edition configuration.

Managing RPVs outside of PowerHA® SystemMirror® control

Once you configure resource groups that contain geographically mirrored volume groups with RPVs, you relinquish control over RPVs to PowerHA® SystemMirror®.

In some cases however, such as when you stop the cluster services, or when you have configured RPVs without yet including geographically mirrored volume groups into the PowerHA® SystemMirror® cluster, you may need to manage the RPV servers and clients manually, outside of the PowerHA® SystemMirror® cluster.

In general, PowerHA® SystemMirror® manages RPVs similar to how it manages volume groups. Before you relinquish control over activation and deactivation of volume groups to PowerHA® SystemMirror®, the groups must be varied off.

Similarly, once RPVs become part of volume groups that belong to an active PowerHA® SystemMirror® cluster, PowerHA® SystemMirror® takes care of changing the RPV states (*defined* and *available*), during cluster events. Before PowerHA® SystemMirror® acquires the RPV servers and clients (by acquiring the resource groups with geographically mirrored volume groups), the RPV servers and clients must be in a *defined* state (that is configured on the system but not active).

In other words, if you stop the cluster services and must manage the RPVs manually, the RPVs must be configured in the *defined* state before you start the cluster services. This enables management of RPVs by PowerHA® SystemMirror® once the nodes join the cluster.

Resuming the RPV clients manually

The RPV client and server are parts of the RPV device driver. An RPV device driver enables the AIX® LVM to consider the physical volume located at the remote site as another local physical volume, although the actual data I/O operations are performed at the remote site.

The following table summarizes the commands that let you recover the RPV client:

| Type of RPV client failure | Command Example |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The RPV client device fails because GLVM for PowerHA® SystemMirror® Enterprise Edition cannot find a valid RPV server address. | To provide a valid RPV server address to the failed RPV client device, use <code>chdev -l hdisk2 -a server_addr=192.168.1.10</code> |
| The RPV client device fails to get a response from the RPV server within the specified I/O timeout period (when an RPV client is managed outside of PowerHA® SystemMirror®). | To request the RPV client device driver to resume I/O operations for the specified RPV client, use <code>chdev -l hdisk2 -a resume=yes</code> |
| The RPV client device is marked as missing by the LVM. | To automatically resume I/O by returning the RPV client device to the available state, use <code>varyonvg <volume_group_name></code> |
| The RPV client device is marked as removed by the LVM. | To automatically resume I/O by returning the RPV client device to the available state, use <code>chpv -v a hdisk2</code> and then use: <code>varyonvg <volume_group_name></code> |

Here is a description of some of these cases in more detail:

Due to certain failures, such as when all XD_data networks become unavailable, or when, for some other reason, the system cannot obtain the RPV server address, the RPV client device can be marked by the system as failed. In this case, use the **chdev** command to provide a valid RPV server address to the RPV client device and to return the RPV client device to the available state and mark it as active, as in the following example:

```
chdev -l hdisk2 -a server_addr=192.168.1.10
```

It is also possible for the RPV client to fail for other reasons that are outside of PowerHA® SystemMirror®. For example, if the RPV client fails to get a response from the RPV server within the specified I/O timeout period, the RPV client fails, even if the RPV server is specified. (This situation can occur if RPVs are used outside of PowerHA® SystemMirror® control).

In this case, to request the RPV client device driver to resume I/O operations for the specified RPV client, use the **chdev** command with the **resume** argument. The value of the resume argument must be **yes**, as in the following example:

```
chdev -l hdisk2 -a resume=yes
```

Managing geographically mirrored volume groups outside of PowerHA® SystemMirror® control

Once volume groups are included in the PowerHA® SystemMirror® cluster resource groups, PowerHA® SystemMirror® takes care of varying the groups on and off, depending on the changing cluster conditions. Before PowerHA® SystemMirror® acquires the volume groups initially, they must be varied off.

Therefore, if for some reason you must stop the cluster services and manage the geographically mirrored volume group outside of the PowerHA® SystemMirror® cluster, use the following tips:

- Before restarting the cluster services, verify that the geographically mirrored volume groups are imported on all the nodes in the cluster on which the associated resource group is configured and on which the volume group is mirrored at the local and remote sites.
- Before restarting the cluster services, verify that the geographically mirrored volume group is varied off.
- If you want to activate the geographically mirrored volume group outside of the PowerHA® SystemMirror® cluster, you may want to verify that the RPV client device for the volume group has a valid RPV server specified for it, otherwise, the RPV client cannot move to the *available* state and the volume group cannot be varied on. In some cases, the RPV client device cannot move to the *available* state due to the fact that the RPV server for it is set to `server_addr=none`.

To provide the RPV client with a valid RPV server IP address, use the `smit rpvclient` fastpath, select **Change/Show Remote Physical Volume Client**, and in the field for **RPV Server IP** address, press F4 to obtain the picklist. Select the valid IP address from the picklist and press Enter.

Troubleshooting GLVM for PowerHA® SystemMirror® Enterprise Edition

These topics present general information for maintaining and troubleshooting a GLVM for PowerHA® SystemMirror® Enterprise Edition configuration. Issues specific to GLVM and RPV support are noted as such.

Logging messages

GLVM for PowerHA® SystemMirror® Enterprise Edition uses the standard logging facilities for PowerHA® SystemMirror®.

To troubleshoot the GLVM for PowerHA® SystemMirror® Enterprise Edition operations in your cluster, use the event summaries in the **hacmp.out** file and the **syslog** file. You can also use the AIX® **trace** command.

Troubleshooting the GLVM Configuration Assistant

If an error occurs while you are using the GLVM Configuration Assistant, the tool exits and displays an error message.

Most of problems that occur when you are using the GLVM Configuration Assistant result from network communication issues between the local site and remote site. Therefore, you must verify that all network communication is working correctly. Other errors might be caused because of a system problem. For example, an AIX® command fails to run correctly.

Any existing GLVM resources in your environment are not affected by using the GLVM Configuration Assistant. However, the GLVM Configuration Assistant might have problems if there are previous definitions of the RPV server and RPV client on the disks that you selected for mirroring.

The removal process that is available in the GLVM Configuration Assistant does not remove the volume group, but removes the mirroring function. The data in the volume group remains on the local node and on the disks that were used by the volume group. You must remove the volume group to delete the data.

Persistent IP labels are IP addresses that are bound to a node and is maintained as a highly available resource by PowerHA® SystemMirror®. For the GLVM Configuration Assistant to work correctly, you must be able to ping the persistent IP labels. You must have as many persistent IP labels as the number of nodes in the site for the cluster. If you have limited persistent IP labels, an error occurs and the GLVM Configuration Assistant exits.

Cache device failure

If the RPV client encounters I/O errors while trying to access the cache logical volume, then it cannot continue to process write requests.

Any cache logical volume I/O failure, even if it is just a single I/O failure, that prevents the RPV driver from being able to perform its duties or to trust the contents of the cache requires the RPV driver to declare that the cache has failed. Then, it must inform LVM that the cache device has failed. LVM keeps track of this by setting a flag that is stored with the other mirror pool attributes in the LVM metadata.

The RPV client must tell LVM to mark ALL of the physical partitions on the disk as stale. Keep in mind that all of the remote physical volumes in the same mirror pool share a single cache device, so a cache device failure affects the entire mirror pool, which means that ALL of the physical partitions in the entire mirror pool can be marked as stale.

The recovery from a cache device failure requires you to perform a set of tasks to remove and replace the cache device. You will also need to run the **syncvg** command to bring the stale physical partitions back up to date. It is necessary for the **syncvg** command to perform a full synchronization across the network. You can protect against cache device failure by using LVM mirroring (two copies of the cache logical volume) or by using a disk subsystem that has built-in data mirroring or RAID capabilities, in order to make the cache logical volume highly available and less likely to fail.

Asynchronous I/O cache logical volume failure

Asynchronous GLVM mirroring requires a **aio_cache** logical volume. A **aio_cache** logical volume stores write requests locally while it waits for the data to be written to a remote disk.

If hardware fails in an asynchronous GLVM configuration, the failure prevents the rpvclient driver from accessing the `aio_cache` logical volume. This type of failure is known as *cache device failure*. When a cache device fails, the **errpt** command displays the following output:

```
# errpt
IDENTIFIER  TIMESTAMP    T C RESOURCE_NAME DESCRIPTION
30097641    0506102109 I S LVDD      AIO CACHE FAIL RECOVERY DONE
EAA3D429    0506102109 U S LVDD      PHYSICAL PARTITION MARKED STALE
BD9355CC    0506102109 U S LVDD      AIO FAIL NOTIFY RECEIVED
F7DDA124    0506102109 U H LVDD      PHYSICAL VOLUME DECLARED MISSING
52715FA5    0506102109 U H LVDD      FAILED TO WRITE VOLUME GROUP STATUS AREA
E86653C3    0506102109 P H LVDD      I/O ERROR DETECTED BY LVM
BD9355CC    0506102109 U S LVDD      AIO FAIL NOTIFY RECEIVED
F7DDA124    0506102109 U H LVDD      PHYSICAL VOLUME DECLARED MISSING
52715FA5    0506102109 U H LVDD      FAILED TO WRITE VOLUME GROUP STATUS AREA
E86653C3    0506102109 P H LVDD      I/O ERROR DETECTED BY LVM
71D2CAA4    0506102109 P S LVDD      AIO CACHE FAIL NOTIFY RECEIVED
```

The `AIO CACHE FAIL NOTIFY RECEIVED` output and the `AIO CACHE FAIL RECOVERY DONE` output indicates that a cache device failed.

The **lsmmp** command also indicates whether that the cache device is valid. When a cache device is not valid, the **lsmmp** command displays the following output:

```
# lsmmp -A datavg
VOLUME GROUP:      datavg          Mirror Pool Super Strict: yes

MIRROR POOL:       siteA           Mirroring Mode:          ASYNC
ASYNC MIRROR STATE: inactive      ASYNC CACHE LV:          datacache2
ASYNC CACHE VALID: yes             ASYNC CACHE EMPTY:       yes
ASYNC CACHE HWM:   80              ASYNC DATA DIVERGED:    no

MIRROR POOL:       siteB           Mirroring Mode:          ASYNC
ASYNC MIRROR STATE: active         ASYNC CACHE LV:          datacache1
ASYNC CACHE VALID: no             ASYNC CACHE EMPTY:       yes
ASYNC CACHE HWM:   100             ASYNC DATA DIVERGED:    no
```

The `ASYNC CACHE VALID` output indicates that the cache device for the mirror pool that is named `siteB` failed.

When a cache device fails, all logical partition copies on the remote disks are marked as stale. To prevent cache device failure, you can use Logical Volume Manager (LVM) mirroring (two copies of the `aio_cache` logical volume) or a disk subsystem that has built-in data mirroring or RAID capabilities that make the `aio_cache` logical volume highly available. Do not put the `aio_cache` logical volume on the same disk as other data to avoid data loss during a disk failure. Data can be lost if local copies and cached updates to remote copies, of the same data are stored on a single disk when the hardware fails. In this scenario, any updates that were written to the failed disk but not mirrored to the remote site are lost.

To fix a cache device failure, complete the following steps:

1. Identify and correct the hardware problem that caused the failure.
2. To switch from asynchronous mirroring to synchronous mirroring, enter the following command from the command line:

```
chmp -S -f -m siteB datavg
```

3. To indicate the remote physical volume (RPV) client can resume communication with the RPV server, for each remote disk, enter the following command:

```
chdev -l hdisk8 -a resume=yes
```

where `hdisk8` is the name of the remote disk.

4. To indicate the LVM to access the disks that were previously marked as missing, enter the following command:

```
varyonvg datavg
```

where `datavg` is the name of the volume group.

You must run the **varyonvg** command when the volume group is in a varied online state.

5. If the hardware failure resulted in the permanent loss of the `aio_cache` logical volume, you must create a new `aio_cache` logical volume.
6. To switch back to asynchronous mirroring, enter the following command:

```
chmp -A -m siteB datavg
```

where `siteB` is the name of the site where the mirror pool exists and `datavg` is the name of the volume group.

7. To verify whether the cache device is valid by identifying whether the `ASYNC CACHE VALID` output is set to yes, enter the following command:

```
lsmp -A datavg
```

where `datavg` is the name of the volume group.

Running corrective actions during verification

When you configure geographically mirrored volume groups, some automatic corrective actions of the **clverify** utility, such as automatically importing or exporting a volume group, do not work. Therefore, run the automatic corrective action of the **clverify** utility in the Interactive mode.

Quorum and forced varyon in GLVM for PowerHA® SystemMirror® Enterprise Edition

Quorum is a feature of the AIX® LVM, and forced varyon for volume groups. How you configure quorum can affect the availability of mirrored data.

Quorum issues

In general, it is recommended to disable quorum for geographically mirrored volume groups in order to minimize the possibility of leaving the resource group in the ERROR state. In PowerHA® SystemMirror®, quorum loss for a volume group causes PowerHA® SystemMirror® to selectively move the affected resource group to another node. Similarly, in GLVM for PowerHA® SystemMirror® Enterprise Edition, if quorum is enabled and then lost for a geographically mirrored volume group, PowerHA® SystemMirror® might start a selective failover event (logged in **hacmp.out** as an **rg_move** for the resource group).

The following statements explain why it is preferable to have a quorum setting disabled. Depending on whether you have quorum enabled or disabled for a geographically mirrored volume group, these scenarios are possible:

- **Quorum is enabled.** If you have configured just one `XD_data` network only and it fails (or, if all of your `XD_data` networks fail), the RPV servers are no longer reachable and the RPV clients fail. This could lead to a quorum loss in some cases. PowerHA® SystemMirror® launches selective failover of the resource group upon the loss of quorum, and attempts to move the resource group after the failure of the `XD_data` network, even though no other node at the site is active on the `XD_data` network. This leaves the resource group in the ERROR state. For information on how to manually recover the resource group, see **No automatic recovery of the resource group in the error state.**
- **Quorum is disabled.** If the `XD_data` network fails (or, if you have more than one `XD_data` network, then all of them fail) and quorum is disabled, it is more likely that access to the active volume group is maintained in the cluster. For instance, in a cluster with two nodes, one at each site, and two disk enclosures at each site, if the `XD_data` network fails, the access to the disks at the remote site is lost. But, if the resource group with the geographically mirrored volume group is hosted at another site (where disks are functioning), it can still remain online.
- Another reason to have quorum disabled is that it allows you to shut down the other site, for maintenance purposes. In this case, half of the disks are not accessible. But since quorum is disabled, and at least one disk remains available at the local site for the resource group, the resource group can still remain online on

the node with access to that disk at the local site. As a result, when you need to perform a site maintenance, PowerHA® SystemMirror® does not attempt to move the resource group to any other node at the remote site.

To summarize, it is recommended to disable quorum for geographically mirrored volume groups and to have more than one data mirroring networks configured. This minimizes unnecessary failovers and also allows GLVM for PowerHA® SystemMirror® Enterprise Edition to use surviving networks for data mirroring and potentially recover the resource group on the node at the remote site.

Note: It is important to realize that having quorum disabled often requires setting forced varyon of the volume group.

Forced varyon issues

You should use a forced varyon option for the geographically mirrored volume groups, along with disabling quorum. Forcing a varyon of volume groups is an option that you should use only with understanding of its consequences.

You can specify a forced varyon attribute for geographically mirrored volume groups, for volume groups on SCSI disks that use LVM mirroring and for volume groups that are mirrored between separate RAID devices.

This section describes the conditions under which you can safely attempt to forcefully bring a geographically mirrored volume group online on the node.

When PowerHA® SystemMirror® attempts a forced varyon

In general, GLVM for PowerHA® SystemMirror® Enterprise Edition attempts a forced varyon in the event of a cluster failure. For troubleshooting purposes, it is helpful to know under what conditions or cluster events GLVM for PowerHA® SystemMirror® Enterprise Edition attempts a forced varyon, when it is configured.

The following list contains examples of cluster event failures that can trigger a forced varyon of a geographically mirrored volume group:

- Cluster startup, normal varyon fails due to a loss of quorum on one of the disks
- Nodes joining the cluster, normal varyon fails due to a loss of quorum on one of the disks
- Selective failover caused by an application or a node failure moves a resource group to a takeover node
- Selective failover caused by a loss of quorum for a volume group moves a resource group to a takeover node.

When a resource group that contains a geographically mirrored volume group falls over to a node at the remote site, GLVM for PowerHA® SystemMirror® Enterprise Edition determines whether any nodes at the remote site have at least one XD_data network up and available and attempts to bring the resource group online on any of these nodes. If no nodes have the XD_data network available, PowerHA® SystemMirror® still attempts to acquire the resource group.

If a resource group is set to forced varyon of volume groups, this may result in a situation when the geographically mirrored volume group is forcefully varied on a node, in the case when the RPV servers are not accessible from the RPV clients (that is, not a single XD_data network is available). If the volume group is forcefully varied on in this case, having an additional network between the sites will prevent the cluster from being partitioned.

Note:

- Be aware that when the forced varyon facility is used successfully and the volume group is brought online on the node (using the one complete copy of the data that was found), the data that you recover by forcing a volume group to go online is guaranteed to be consistent, but not necessarily the latest.
- During runtime, for large volume groups (those with more than 256 disks), checking logical partition maps may take extra processing time. However, enduring a slow varyon process that enables data recovery is preferable to having no chance at all to activate the volume group.

Avoiding a partitioned cluster

The forced option to activate a geographically mirrored volume group must be used with care. Should the cluster become partitioned, each partition might force on the volume group, and continue to run. In this case, two unequal copies of the data will be active at the same time. This situation can cause data divergence and does not allow a clean recovery.

To prevent cluster partitioning, configure multiple networks between the sites.

When it is no longer possible for the two sites to communicate using any of the configured XD_data networks, but there is still a XD_ip network available for heartbeating, GLVM for PowerHA® SystemMirror® Enterprise Edition issues a **site_isolation** event.

In this case, the geographically mirrored volume group continues to be activated on a node at the local site (that is, the primary instance of the resource group remains online), but the mirroring is stopped. After you restore at least one of your previously configured XD_data networks and the cluster is stable, the mirroring resumes automatically.

Verification checks for forced varyon

If you specified a forced varyon attribute for a volume group in the resource group, and PowerHA® SystemMirror® detects that the logical volumes are not being mirrored with the super strict disk allocation policy, it issues a warning during verification of cluster resources. (GLVM for PowerHA® SystemMirror® Enterprise Edition prevents you from configuring geographically mirrored volume groups if the disk allocation policy is not super strict).

In addition, PowerHA® SystemMirror® checks the logical partitions on each disk for each volume group:

- If it cannot find a complete copy of every logical volume for a volume group, an error message: "Unable to vary on volume group <vg name> because logical volume <logical volume name> is incomplete" displays in the **hacmp.out** file. In this case, a forced varyon operation fails and you will see an event error.
- If PowerHA® SystemMirror® can find a complete copy for every logical volume for all volume groups in this resource group that require a forced varyon, the volume groups get varied on the node in the cluster.

Note: In addition to this check done by PowerHA® SystemMirror®, an additional check is being done by GLVM for PowerHA® SystemMirror® Enterprise Edition. In some cases, even if a complete copy of every logical volume is found, the software might still determine to not forcefully activate the volume group.

GLVM for PowerHA® SystemMirror® Enterprise Edition does not forcefully activate a volume group

In some instances the software detects that the resource group is about to access the stale partitions on the mirror copy. This is related to a situation when, after a failover to another site, GLVM for PowerHA® SystemMirror® Enterprise Edition determines that the most recent mirroring updates were possibly not done.

In these cases, GLVM for PowerHA® SystemMirror® Enterprise Edition does not forcefully activate the volume group (even if forced varyon is set to True). The information about it is displayed in the **hacmp.out** file, as described in the following example.

The following example explains when GLVM for PowerHA® SystemMirror® Enterprise Edition skips the forced varyon process and does not forcefully activate a volume group:

Typically, when resource group recovery is needed, GLVM for PowerHA® SystemMirror® Enterprise Edition performs a failover and moves the resource group from its local site to a node at another site. Before activating the volume group, it checks whether the mirror copies are updated with the most recent updates, including those that would be done just before PowerHA® SystemMirror® moves the active resource group from its local node to another site. If it determines that the mirror copy on the remote site may not contain the most recent updates, the volume group is not forced to varyon at the remote site, sending the resource group into the ERROR state on the node at the remote site.

If the resource group is in ERROR state, **hacmp.out** shows that the volume group failed to vary on.

Note: The SKIP_FORCED_VARYON flag can be set both in cases when quorum is enabled and when it is disabled, if GLVM for PowerHA® SystemMirror® Enterprise Edition determines that mirroring updates were not done before a failover occurred. The following example describes a case with quorum enabled.

The entries in **hacmp.out** are as follows:

```
rg7:glvm_predisk_available[529] [[ -z '' ]]  
rg7:glvm_predisk_available[531] [[ -n duke ]]  
rg7:glvm_predisk_available[533] STATUS=4  
rg7:glvm_predisk_available[538] [[ NONE == RELEASE_SECONDARY ]]  
rg7:glvm_predisk_available[554] exit 4  
rg7:get_disk_vg_fs[192] export SKIP_FORCED_VARYON=true  
.  
.  
.  
rg7:clvaryonvg[504] varyonvg -n gmv  
0516-052 varyonvg: Volume group cannot be varied on without a quorum.  
More physical volumes in the group must be active. Run diagnostics on  
inactive PVs.  
.  
.  
.  
rg7:cl_mirrorset[87] [[ -n true ]]  
rg7:cl_mirrorset[89] return 1
```

Clearing a skip_forced_varyon setting and activating a volume group

To summarize, the event processing in GLVM for PowerHA® SystemMirror® Enterprise Edition does not force a varyon on a volume group if PowerHA® SystemMirror® has a reason to believe that the disks used by the volume group may contain stale data. If PowerHA® SystemMirror® decides this is the case, it sets the flag SKIP_FORCED_VARYON=true.

In this case, you can decide if not varying on the volume group is indeed the proper thing to do based on the conditions in the cluster. For instance, you may want to use this data if you are completely sure that the other site has failed and is unrecoverable, because stale data is better than no data at all. However, you do not want to use this copy of the data after a planned site outage, when the latest copy of the data is still intact at the other site.

Carefully investigate the conditions in the cluster and if you are confident about the integrity of the data in the volume group, you can manually activate the resource group on the node. (This will set the SKIP_FORCED_VARYON=false). For instance, manually activating the resource group could be useful in a disaster recovery situation when activating a copy even if it is not the latest is preferable to not having any copies recovered.

To forcefully activate a volume group in this case, manually bring the resource group online on the node on which the resource group went into the ERROR state.

Data availability and data integrity

There are several general types of cluster configurations that are possible in GLVM for PowerHA® SystemMirror® Enterprise Edition.

Quorum disabled and forced varyon set to true

If quorum is disabled for a volume group and the volume group is allowed to be varied on forcefully, high availability of your data is achieved, but in rare cases data integrity can be at risk. This section describes the scenario and explains how to prevent it.

Note: Quorum disabled and forced varyon set to **True** is a recommended configuration for GLVM for PowerHA® SystemMirror® Enterprise Edition.

The situation is possible in which a resource group accesses a copy of the disk where data is not synchronized with the copy on the remote site. This scenario can occur under the following conditions:

- All configured XD_data networks are down (if you have one XD_data, then it is down) and the mirroring had stopped. At this time, due to a cluster condition, the resource group falls over to another site and is activated with a volume group forcefully varied on at a node at a remote site (in other words, secondary instance of the resource group is online at the remote site).
If, under these conditions the node at the local site is restarted and activates the resource group, data is not synchronized.

Preventing data divergence

- In general, if you know that mirroring was not happening the last time before the failure of the last remaining XD_data network (if you had more than one XD_data configured), do not restart the resource group on another site. You may still have a copy of the volume group on the local site, although the data may not be the latest.
- Do not start the node at the local site if all XD_data networks are down, the resource group is online on the node at a remote site (with a forcefully varied on volume group) and you know that prior to the network failure the resource group fell over to the remote site and was online on a node at the remote site.
- Do not stop the cluster services on the higher priority node for the resource group at the local site if all XD_data networks are down and the resource group is activated with a volume group forcefully varied on at a node at a remote site. You do not want to create conditions under which a failover will occur to the site that initially was the local site for this resource group.
- Do not restart the cluster services on the higher priority node.

You can consider manually recovering the updates to the volume group that were done during the time that the network was down and the resource group remained at the remote site.

Quorum enabled and forced varyon set to false

If you do not have the forced varyon option set for the volume group in the resource group but quorum is enabled, the resource group goes into the ERROR state after PowerHA® SystemMirror® attempts to activate it on all available nodes at the remote site.

In this case, while your cluster is protected from data divergence and data integrity is secure, data availability is not always achievable automatically.

If you have this combination of settings, then in certain cases when it is completely safe to vary on the volume group, GLVM for PowerHA® SystemMirror® Enterprise Edition moves the resource group into an ERROR state, instead of attempting to activate it on a node where access to a single copy of the data exists.

Activating the volume group when it is safe

- To ensure that you are aware of this situation and can take action to activate the volume group, do the following:

- Create a post-event script or an error notification method to take action or notify you when the resource group goes into an ERROR state.
- Manually take the resource group offline on the node where it is in the ERROR state using the PowerHA® SystemMirror® Resource Group and Application Management utility in SMIT (clRGmove utility). Set forced varyon to True for the volume group and manually bring the resource group online on the node.

Shutting down and restarting the cluster services at both sites

In some cases when the last of all configured XD_data networks goes down but you know that no data is lost, you may want to consider shutting down the cluster services on both sites to restore the network.

This could happen if a resource group fell over to another site and is in the ERROR state, and all configured XD_data networks are down. If you shut down the cluster services on both sites, it is very important to carefully investigate on which node the resource group was last active and mirrored, and to display the cluster services in the correct order on the nodes. Ensure that the order in which you restart the cluster services for the sites is correct: Start the cluster services *first* on the nodes at the site at which the latest mirrored instance of the resource group was active.

Selective fallover

As in base PowerHA® SystemMirror®, GLVM for PowerHA® SystemMirror® Enterprise Edition supports selective fallover for a resource group that contains a geographically mirrored volume group, and attempts to move the resource group to another node.

If all XD_data networks fail, and if the node on which the network failed is hosting a geographically mirrored volume group that belongs to the primary copy of the resource group, PowerHA® SystemMirror® checks if there are any other nodes available at the local site that have an available IP interface on any of the XD_data networks. If such a node exists, PowerHA® SystemMirror® selectively moves the resource group to this node. Otherwise, the group remains online on its current owner node (but the mirroring to the remote site stops). After at least one of the previously configured XD_data networks is restored on the node, the mirroring function is resumed automatically and the copies are synchronized.

Note that PowerHA® SystemMirror® only applies selective fallover for primary copies of the resource groups, and not for the resource group instances accessing copies of the data on the remote site. That is, if the last remaining XD_data network fails, and the node on which the network failed is hosting a resource group that accesses a mirror copy of the volume group through associated RPV servers (this is also referred to as a secondary instance of the resource group), GLVM for PowerHA® SystemMirror® Enterprise Edition does not move the resource group to another node at the remote site and the resource group goes into the ONLINE_SECONDARY state. The resource group remains in the ONLINE_SECONDARY state and does not participate in mirroring. After the XD_data network is restored, the mirroring function is resumed automatically.

Resource group recovery

When a resource group that contains a geographically mirrored volume group falls over to another node or you move it to another node within the same site using SMIT, GLVM for PowerHA® SystemMirror® Enterprise Edition determines whether any nodes at this site have an XD_data network up and available and attempts to bring the resource group online on any of these nodes. If no nodes have the XD_data network available, PowerHA® SystemMirror® still attempts to acquire the resource group.

If forced varyon is set to True, this may result in a situation when the geographically mirrored volume group is forcefully varied on a node, in the case when the RPV servers are not accessible from the RPV clients and the mirroring function will not be working. In this case, when the XD_data network is restored, manual intervention is required to reestablish the mirroring operations.

Resource group fallover from the local to the remote site

GLVM for PowerHA® SystemMirror® Enterprise Edition attempts to move the resource group from the local to the remote site only when no nodes at the local site can support the resource group in any form. If the node where the resource group is ONLINE fails and there is no takeover node available at the same site, GLVM for PowerHA® SystemMirror® Enterprise Edition attempts to recover the resource group on an available takeover node at the remote site.

No automatic recovery of the resource group in the error state

GLVM for PowerHA® SystemMirror® Enterprise Edition does not automatically attempt to recover a resource group that contains a geographically mirrored volume group when the resource goes into the ERROR state.

For instance, if the resource group that contains a geographically mirrored volume group is in the ERROR state and all XD_data networks are down, then if one of the XD_data networks comes back up PowerHA® SystemMirror® does not attempt to recover the resource group and bring it ONLINE. You must manually bring the resource group online on a node that has an active XD_data network connection.

Before reactivating the resource group that went into the ERROR state and varying on the geographically mirrored volume group, wait until all local and remote copies of the data are available again.

If you decide to bring online the resource group with only the local or remote copy of the data available, be aware that you may not have the latest copy of the data, and that you may have two diverged copies of the data.

Recovering from a disaster while using asynchronous mirroring

If the production site suddenly fails before the volume group can be varied offline, the physical volumes at the disaster recovery site are likely to be missing some or all of the updates that are still stored in the cache at the production site.

If this situation occurs, you have two options. You can either wait until the production site is recovered, if that is possible, in which case you will not lose the data updates that are stored in the cache. Or you can activate the volume group at the disaster recovery site without the latest data updates, in which case you are likely to end up with data divergence, if the data at the production site has not been destroyed.

Use the **varyonvg** command to bring the volume group online at the disaster recovery site. A normal **varyonvg** command usually fails because the production site disks are not accessible. When quorum is enabled, more than half of the disks must be accessible for normal **varyonvg** to succeed. When quorum is disabled, all of the disks must be accessible for normal **varyonvg** to succeed. You should consider disabling quorum, so normal **varyonvg** should fail in the vast majority of cases. For this reason, a forced **varyonvg** is normally required in order to bring the volume group online after a production site failure. However, this fails when asynchronous mirroring is being used, to prevent data divergence from occurring by mistake.

In this situation, you must decide whether or not to allow data divergence to occur. If you chose to run your applications at the disaster recovery site, you need to run the **varyonvg** command with the **-d** flag (combined with the **-f** flag) to allow the volume group to be brought online at the disaster recovery site without the latest updates that were lost when the production site failed. In the process of doing this, the LVM records the fact that data divergence has been allowed to occur.

Recovering the production site

The production site can fail for a variety of reasons. An event such as a fire or flood in the data center can destroy the hardware. In that case, it is not possible to recover the production site and the data that was stored there is lost. However, a system crash or power outage can cause a production site outage without destroying the data that resides there.

If the volume group was brought online at the disaster recovery site without the latest updates, and if the data on the production site disks is still intact, then production site recovery requires you to decide what to do about data divergence. You need to decide which copy of data to keep and which copy should be discarded.

Recovering from data divergence

Data divergence is a state where each site's disks contain data updates that have not been mirrored to the other site.

Each site's copy of the data reflects logical volume writes that are missing from the other site's copy of the data. If data divergence is allowed to occur, you must later choose to keep the data that exists at one site and back out, in other words throw away, the non-mirrored updates that occurred at the other site, before the entire volume group can be merged back together.

Recover from data divergence

When you are recovering your production site and data divergence has occurred, you can choose to allow data divergence recovery.

If you do want to allow data divergence recovery to occur, then you must indicate whether to preserve the data copy that resides at the "local" or "remote" site. In this case "local" refers to the local site, where you are running the recovery, and "remote" refers to the opposite site. The **varyonvg** command processing preserves the data at the site you select and backs out the non-mirrored updates at the opposite site, effectively throwing them away, in order to merge the volume group back together. Note that LVM allows up to three copies of each logical volume. Therefore, one site may actually contain two copies of your data. However, during recovery, you are merely indicating which site has the data that you want to keep. If the site that you select has two mirror copies, then both mirror copies will be preserved.

Deciding which copy to preserve often depends on which site has more non-mirrored updates. If your applications ran for several hours at the disaster recovery site, then you may want to keep the data that resides at the disaster recovery site, and lose the non-mirrored updates that reside in the cache at the production site. On the other hand, if you did not make any meaningful changes to your data at the disaster recovery site--perhaps you only mounted your file systems--then you may want to keep the data at the production site. In this case, the data divergence recovery processing would back out the data updates that were made during the mounting of the file systems. The decision of which copy of the data to keep must be made based on the circumstances, and only you are able to make this decision.

In certain cases you cannot choose to keep the data at the production site. For example, if you removed logical volumes and created new logical volumes while running at the disaster recovery site, while the production site was down, then the **varyonvg** command may determine that it cannot safely back out the non-mirrored updates from the disaster recovery site copy, and in that case the **varyonvg** command will fail.

Forcing data divergence recovery

Whichever data copy you choose to keep effectively becomes the "latest" copy of your data. If the site that you select has only one mirror copy of the data, and that copy contains stale physical volumes, then the **varyonvg** command will fail. (Likewise, if the site you select has two mirror copies of the data and both copies of a given logical partition are stale, in other words both physical partitions for a given logical partition are stale, then the **varyonvg** command will fail.) If the data had not diverged, then a stale physical partition could be synchronized from the data at the opposite site. However, the opposite site contains data that you have chosen to discard. The **varyonvg** command does not understand your data in order to determine whether it is safe to do this without the risk of compromising the integrity and consistency of your data.

If the **varyonvg** command fails because it finds stale physical partitions (or both logical partition copies are stale if there are two mirror copies) in the data version that you intend to keep, then you have one of two choices. You can decide to keep the data at the opposite site, if that is appropriate. For example, if the production site failed while it was running the **syncvg** command to update the disks at the disaster recovery site, then **varyonvg** will fail when it finds stale physical partitions on the disaster recovery site physical volumes that had not been updated by the **syncvg** command. If you are willing to lose whatever changes were made while the volume group was varied online at the disaster recovery site, then you may want to keep the production site version of the data instead of the disaster recovery site version. Then, the next time that the **syncvg** command runs, it can finish synchronizing the data at the disaster recovery site.

On the other hand, if you know that it is safe to do so, you may choose to force the **varyonvg** command to keep the data copy that contains stale physical partitions. In that case, the stale physical partitions will be synchronized from the data copy that resides at the opposite site, from the data that is about to be discarded. **HOWEVER, YOU ASSUME THE RESPONSIBILITY FOR THE INTEGRITY OF YOUR DATA IF YOU USE THIS OPTION.** When this option is used, the **varyonvg** command prints a warning message to identify which logical volumes were recovered in this manner. Therefore, this option may be useful if you have a large number of logical volumes and only a small subset have stale physical partitions. In that case, the logical volumes which did not have warning messages were recovered safely, and the logical volumes that did have warning messages may have been compromised. Then, you may choose to recover those logical volumes by restoring from a backup image. This option may be more desirable than restoring all of your logical volumes from backup images.

Overriding default data divergence recovery

When recovering from a production site failure, you must decide how to handle data divergence recovery before the production site can be successfully brought back into the cluster.

If you have a default selection for handling data divergence for the resource group, you can choose to have PowerHA® SystemMirror® automatically perform data divergency recovery. Event processing will automatically handle data divergency, regardless of whether the resource group is processed automatically while starting cluster services or if you manage the resource group manually.

You need to bring the secondary instance of the resource group online before the primary instance is brought online, if the secondary instance of the resource group is the one that contains the data to be preserved. For example, if you want to keep the data at the disaster recovery site, then you should not take the resources offline at the disaster recovery site prior to bringing them online at the production site. Otherwise, the **varyonvg** command will not have access to the disks at both sites in order to detect that data divergence has occurred, and the resource group will be brought online with the wrong version of the data. Do not shutdown or take resources offline at the disaster recovery site, if it contains the data to be preserved, prior to starting up PowerHA® SystemMirror® on the production site.

If you need to override a resource group's default data divergence recovery attribute, then it is necessary to start cluster services on the production site without allowing automatic resource group management. Then you must manage the resource group manually, and at that time you can specify a different data divergence recovery value that will override the default value.

First, start cluster services without allowing the joining nodes to acquire any resources. This is accomplished by selecting **Manually** for the Manage Resource Groups field, as shown here:

```

Start Cluster Services

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

* Start now, on system restart or both
  Start Cluster Services on these nodes
* Manage Resource Groups
  BROADCAST message at startup?
  Startup Cluster Information Daemon?
  Ignore verification errors?
  Automatically correct errors found during
  cluster start?

                                [Entry Fields]
                                now                +
                                [nodeA1]           +
                                Manually           +
                                false              +
                                false              +
                                false              +
                                Interactively      +

```

Next, after the production site nodes joins the cluster, you must manage the resource groups manually. Most likely the primary instance of the resource group is already running at the disaster recovery site. You have a few choices:

- Move the primary instance of the resource group back to the production site. This performs what is known as a “site fallback” to return the cluster to how it was prior to the site failure, with the production site asynchronously mirroring to the disaster recovery site. PowerHA® SystemMirror® automatically takes the primary resource group instance offline and then bring the secondary instance online at the disaster recovery site before bringing the primary instance online at the production site.
- Keep the primary instance at the disaster recovery site, and bring the secondary instance of the resource group online at the production site. Then asynchronous data mirroring occurs in the opposite direction, from the disaster recovery site back to the production site.
- If you want to switch back to the production site's version of the data, while continuing to run at the disaster recovery site, then it will be necessary to take the resource group offline at the disaster recovery site. Then you can bring up the secondary instance at the production site, followed by the primary instance at the disaster recovery site.

You can specify which site's version of the data should be preserved by data divergence recovery processing on the **Bring a Resource Group Online and Move a Resource Group to Another Node / Site** SMIT menu, as shown here:

```

Bring a Resource Group Online

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

Resource Group to Bring Online
Node on Which to Bring Resource Group Online
Name of site containing data to be preserved
  by data divergence recovery processing
  (Asynchronous GLVM Mirroring Only)

                                [Entry Fields]
                                glvm_rg
                                nodeA1
                                [siteA]          +

```

You can either fill in the **Name of site containing data to be preserved by data divergence recovery processing** field with a site name or leave it blank. If you leave the field blank, then the value that was previously specified is used for event process. If there isn't a value specified, then this operation fails.

If a site name is specified, the name is passed to event processing for the resource group for this operation only.

Monitoring the state of your GLVM environment

To assist you in monitoring the state of your GLVM environment, including the RPVs and GMVGs you have configured, GLVM includes two tools: `rpvstat` and `gmvstat`. These two commands provide real time status information about RPVs and GMVGs.

The `rpvstat` command

The `rpvstat` command provides status monitoring for RPV Clients.

It displays the following information for one or more RPV Clients:

- RPV client name
- Connection status
- Total number of completed reads
- Total number of KBs read
- Total number of read errors
- Total number of pending reads
- Total number of pending KBs to read
- Total number of completed writes
- Total number of KBs written
- Total number of write errors
- Total number of pending writes
- Total number of pending KBs to write

The `rpvstat` command can optionally display its I/O-related statistics on a per-network basis. The network summary option displays the following additional information:

- Network throughput in KBs per second

The `rpvstat` command can also display the highest recorded values for the pending statistics. These historical high water mark numbers are:

- Maximum number of pending reads per device and network (high water mark)
- Maximum number of pending KBs to read per device and network (high water mark)
- Maximum number of pending writes per device and network (high water mark)
- Maximum number of pending KBs to write per device and network (high water mark)
These statistics are reported on a separate display and include the additional statistic:
- Number of retried I/O operations (combination of both reads and writes)

The `rpvstat` command allows information to be displayed for all RPV clients on the system or for a subset of RPV clients specified by RPV client name on the command line. The `rpvstat` command also allows the information to be monitored (redisplayed at a user-specified interval).

The `rpvstat` command interacts with the RPV client pseudo device driver to retrieve the information the client displays.

Sample display output for the rpvstat command

These samples display various output for the rpvstat command.

Example 1

Run the **rpvstat** command with no options to display each RPV client on the local node, along with accumulated statistics.

```
>>>rpvstat
```

| RPV Client | cx | Comp Reads Pend Reads | Comp Writes Pend Writes | Comp KRead Pend KRead | Comp KWrite Pend KWrite | Errors |
|------------|----|--------------------------|----------------------------|--------------------------|----------------------------|--------|
| hdisk144 | 4 | 22482004 43339 | 0 0 | 22834857 3971740 | 0 0 | 384 |
| hdisk158 | 3 | 488700 10 | 0 0 | 8336500 5120 | 0 0 | 122 |
| hdisk173 | 0 | 148839202 73980 | 0 0 | 294827470 6349820 | 0 0 | 539 |
| hdisk202 | X | 0 0 | 0 0 | 0 0 | 0 0 | 0 |

Example 2

Run the **rpvstat** command with no options, but specify a single RPV client to display accumulated statistics for that particular RPV client only.

```
>>>rpvstat hdisk158
```

| RPV Client | cx | Comp Reads Pend Reads | Comp Writes Pend Writes | Comp KRead Pend KRead | Comp KWrite Pend KWrite | Errors |
|------------|----|--------------------------|----------------------------|--------------------------|----------------------------|--------|
| hdisk158 | 4 | 488700 5120 | 0 0 | 8336500 122 | 0 0 | 122 |

Example 3

Run the **rpvstat** command with the **-n** option to additionally show accumulated RPV client statistics for each currently defined network.

```
>>>rpvstat -n
```

| RPV Client | cx | Comp Reads Pend Reads | Comp Writes Pend Writes | Comp KRead Pend KRead | Comp KWrite Pend KWrite | Errors |
|----------------|----|--------------------------|----------------------------|--------------------------|----------------------------|--------|
| hdisk144 | 4 | 22482004 339 | 0 0 | 22834857 71740 | 0 0 | 384 |
| 103.17.133.102 | Y | 5620504 81 | 0 0 | 5100384 18038 | 0 0 | 81 |
| 103.17.133.104 | Y | 5620500 85 | 0 0 | 4892247 20448 | 0 0 | 101 |
| 103.17.133.202 | Y | 5620598 85 | 0 0 | 5822041 16384 | 0 0 | 98 |
| 103.17.133.204 | Y | 5620502 84 | 0 0 | 7020185 16870 | 0 0 | 104 |
| hdisk158 | 3 | 488700 10 | 0 0 | 8336500 5120 | 0 0 | 122 |
| 103.17.133.102 | Y | 122175 2 | 0 0 | 2084100 1280 | 0 0 | 28 |
| 103.17.133.104 | N | 122174 2 | 0 0 | 2000311 1288 | 0 0 | 32 |
| 103.17.133.202 | Y | 122176 3 | 0 0 | 2118732 1284 | 0 0 | 30 |
| 103.17.133.204 | Y | 122175 3 | 0 0 | 2133357 1268 | 0 0 | 32 |
| hdisk173 | 0 | 148839202 73980 | 0 0 | 294827470 6349820 | 0 0 | 539 |

| | | | | | | |
|----------------|---|-------------------|--------|---------------------|--------|-----|
| 103.17.133.102 | N | 37209801 18495 | 0 0 | 73706868 1239876 | 0 0 | 100 |
| 103.17.133.104 | N | 37209801 18495 | 0 0 | 73706867 1843755 | 0 0 | 80 |
| 103.17.133.202 | N | 37209800 18495 | 0 0 | 73706868 1533984 | 0 0 | 150 |
| 103.17.133.204 | N | 37209800 18495 | 0 0 | 73706867 1732205 | 0 0 | 209 |
| hdisk202 | X | 0 0 | 0 0 | 0 0 | 0 0 | 0 |

Example 4

Run the **rpvstat** command specifying a single RPV client, a monitor interval of 30 seconds and 4 repeats, and a display of the date and time for each interval. When running in monitor mode with the **-d** option, some of the repeated statistics show only the delta from their previous value, as indicated by the preceding plus sign (+).

```
>>>rpvstat -t -i 30 -c 4 -d hdisk158
```

```
Remote Physical Volume Statistics                               Thu Jan 4 18:11:35 2007
```

| RPV Client | cx | Comp Reads Pend Reads | Comp Writes Pend Writes | Comp KRead Pend KRead | Comp KWrite Pend KWrite | Errors |
|------------|----|--------------------------|----------------------------|--------------------------|----------------------------|--------|
| hdisk158 | 3 | 488700 10 | 0 0 | 8336500 5120 | 0 0 | 122 |

```
Remote Physical Volume Statistics                               Thu Jan 4 18:12:05 2007
```

| RPV Client | cx | Comp Reads Pend Reads | Comp Writes Pend Writes | Comp KRead Pend KRead | Comp KWrite Pend KWrite | Errors |
|------------|----|--------------------------|----------------------------|--------------------------|----------------------------|--------|
| hdisk158 | 3 | +10 8 | +0 0 | +5120 4096 | +0 0 | +0 |

```
Remote Physical Volume Statistics                               Thu Jan 4 18:12:35 2007
```

| RPV Client | cx | Comp Reads Pend Reads | Comp Writes Pend Writes | Comp KRead Pend KRead | Comp KWrite Pend KWrite | Errors |
|------------|----|--------------------------|----------------------------|--------------------------|----------------------------|--------|
| hdisk158 | 3 | +100 4 | +1 0 | +51200 2048 | +0 0 | +0 |

```
Remote Physical Volume Statistics                               Thu Jan 4 18:13:05 2007
```

| RPV Client | cx | Comp Reads Pend Reads | Comp Writes Pend Writes | Comp KRead Pend KRead | Comp KWrite Pend KWrite | Errors |
|------------|----|--------------------------|----------------------------|--------------------------|----------------------------|--------|
| hdisk158 | 3 | +182 11 | +0 0 | +56320 5632 | +0 0 | +0 |

Example 5

Run the **rpvstat** command with the **-n** flag to show individual statistics for each network, specifying a single RPV client, and monitoring every 30 seconds for 4 repeats.

```
>>>rpvstat -n -i 30 -c 4 -d hdisk158
```

```
Remote Physical Volume Statistics
```

| RPV Client | cx | Comp Reads Pend Reads | Comp Writes Pend Writes | Comp KRead Pend KRead | Comp KWrite Pend KWrite | Errors |
|----------------|----|--------------------------|----------------------------|--------------------------|----------------------------|--------|
| hdisk158 | 4 | 488700 10 | 0 0 | 8336500 5120 | 0 0 | 122 |
| 103.17.133.102 | Y | 122175 2 | 0 0 | 2084100 1280 | 0 0 | 28 |
| 103.17.133.104 | Y | 122174 2 | 0 0 | 2000311 1288 | 0 0 | 32 |
| 103.17.133.202 | Y | 122176 3 | 0 0 | 2118732 1284 | 0 0 | 30 |
| 103.17.133.204 | Y | 122175 | 0 | 2133357 | 0 | 32 |

| | | 3 | 0 | 1268 | 0 | | | |
|-----------------------------------|----|--------------------|-------|---------------------|--------|--------------------|-------------------|--------|
| Remote Physical Volume Statistics | | | | | | | | |
| RPV Client | cx | Comp Pend Reads | Reads | Comp Pend Writes | Writes | Comp Pend KByte | KBRead KBWrite | Errors |
| hdisk158 | 4 | +10 | | +0 | | +5120 | +0 | +0 |
| | | 8 | | 0 | | 4096 | 0 | |
| 103.17.133.102 | Y | +3 | | +0 | | +1200 | +0 | +0 |
| | | 2 | | 0 | | 1024 | 0 | |
| 103.17.133.104 | Y | +3 | | +0 | | +1210 | +0 | +0 |
| | | 2 | | 0 | | 1024 | 0 | |
| 103.17.133.202 | Y | +2 | | +0 | | +1350 | +0 | +0 |
| | | 3 | | 0 | | 1024 | 0 | |
| 103.17.133.204 | Y | +5 | | +0 | | +1360 | +0 | +0 |
| | | 3 | | 0 | | 1024 | 0 | |
| Remote Physical Volume Statistics | | | | | | | | |
| RPV Client | cx | Comp Pend Reads | Reads | Comp Pend Writes | Writes | Comp Pend KByte | KBRead KBWrite | Errors |
| hdisk158 | 4 | +100 | | +0 | | +51200 | +0 | +1 |
| | | 4 | | 0 | | 2048 | 0 | |
| 103.17.133.102 | Y | +25 | | +0 | | +12000 | +0 | +0 |
| | | 2 | | 0 | | 512 | 0 | |
| 103.17.133.104 | Y | +25 | | +0 | | +13500 | +0 | +0 |
| | | 2 | | 0 | | 512 | 0 | |
| 103.17.133.202 | Y | +25 | | +0 | | +13600 | +0 | +1 |
| | | 3 | | 0 | | 512 | 0 | |
| 103.17.133.204 | Y | +25 | | +0 | | +12100 | +0 | +0 |
| | | 3 | | 0 | | 512 | 0 | |
| Remote Physical Volume Statistics | | | | | | | | |
| RPV Client | cx | Comp Pend Reads | Reads | Comp Pend Writes | Writes | Comp Pend KByte | KBRead KBWrite | Errors |
| hdisk158 | 4 | +182 | | +0 | | +56320 | +0 | +0 |
| | | 11 | | 0 | | 5632 | 0 | |
| 103.17.133.102 | Y | +46 | | +0 | | +11184 | +0 | +0 |
| | | 2 | | 0 | | 1245 | 0 | |
| 103.17.133.104 | Y | +45 | | +0 | | +18442 | +0 | +0 |
| | | 2 | | 0 | | 1530 | 0 | |
| 103.17.133.202 | Y | +45 | | +0 | | +9942 | +0 | +0 |
| | | 3 | | 0 | | 1302 | 0 | |
| 103.17.133.204 | Y | +46 | | +0 | | +16752 | +0 | +0 |
| | | 3 | | 0 | | 1555 | 0 | |

Example 6

Run the **rpvstat** command with the **-N** option to display summary statistics for each mirroring network. Monitor every 30 seconds for a total of 4 repeats. This can be used to detect errors on a particular network.

```
>>>rpvstat -N -i 30 -c 4 -d
```

| Remote Physical Volume Statistics | | | | | | | |
|-----------------------------------|--------------------|-------|---------------------|--------|--------------------|-------------------|------------------|
| RPV Client Network | Comp Pend Reads | Reads | Comp Pend Writes | Writes | Comp Pend KByte | KBRead KBWrite | Errors KB/sec |
| 103.17.133.102 | 122175 | | 0 | | 2084100 | 0 | 28 |
| | 2 | | 0 | | 1280 | 0 | - |
| 103.17.133.104 | 122174 | | 0 | | 2000311 | 0 | 32 |
| | 2 | | 0 | | 1288 | 0 | - |
| 103.17.133.202 | 122176 | | 0 | | 2118732 | 0 | 30 |
| | 3 | | 0 | | 1284 | 0 | - |
| 103.17.133.204 | 122175 | | 0 | | 2133357 | 0 | 32 |
| | 3 | | 0 | | 1268 | 0 | - |
| Remote Physical Volume Statistics | | | | | | | |
| RPV Client Network | Comp Pend Reads | Reads | Comp Pend Writes | Writes | Comp Pend KByte | KBRead KBWrite | Errors KB/sec |
| 103.17.133.102 | +100 | | +0 | | +123 | +0 | +0 |
| | 7 | | 0 | | 20 | 0 | 4 |
| 103.17.133.104 | +22 | | +0 | | +58 | +0 | +0 |

| | | | | | |
|-----------------------------------|--------------------------|----------------------------|--------------------------|----------------------------|------------------|
| | 2 | 0 | 13 | 0 | 2 |
| 103.17.133.202 | +4 | +0 | +101 | +0 | +1 |
| | 11 | 0 | 82 | 0 | 3 |
| 103.17.133.204 | +88 | +0 | +28 | +0 | +0 |
| | 9 | 0 | 2 | 0 | 1 |
| Remote Physical Volume Statistics | | | | | |
| RPV Client Network | Comp Reads Pend Reads | Comp Writes Pend Writes | Comp KRead Pend KRead | Comp KWrite Pend KWrite | Errors KB/sec |
| ----- | ----- | ----- | ----- | ----- | ----- |
| 103.17.133.102 | +0 | +0 | +0 | +0 | +0 |
| | 0 | 0 | 0 | 0 | 0 |
| 103.17.133.104 | +2 | +0 | +123 | +0 | +0 |
| | 0 | 0 | 0 | 0 | 4 |
| 103.17.133.202 | +0 | +0 | +0 | +0 | +0 |
| | 1 | 0 | 0 | 0 | 0 |
| 103.17.133.204 | +10 | +0 | +28 | +0 | +0 |
| | 0 | 0 | 17 | 0 | 1 |
| Remote Physical Volume Statistics | | | | | |
| RPV Client Network | Comp Reads Pend Reads | Comp Writes Pend Writes | Comp KRead Pend KRead | Comp KWrite Pend KWrite | Errors KB/sec |
| ----- | ----- | ----- | ----- | ----- | ----- |
| 103.17.133.102 | +100 | +0 | +28 | +0 | +1 |
| | 7 | 0 | 21 | 0 | 1 |
| 103.17.133.104 | +22 | +0 | +12 | +0 | +0 |
| | 2 | 0 | 13 | 0 | 0 |
| 103.17.133.202 | +4 | +0 | +8 | +0 | +0 |
| | 11 | 0 | 8 | 0 | 0 |
| 103.17.133.204 | +88 | +0 | +374 | +0 | +0 |
| | 9 | 0 | 123 | 0 | 12 |

Example 7

Run the **rpvstat** command with the **-m** option to display maximum pending statistics (high water marks). This displays the high water mark statistics first by RPV device (for all networks) and then by network (for all devices).

```
>>>rpvstat -m
```

| RPV Client | cx | Maximum Pend Reads | Maximum Pend Writes | Maximum Pend KRead | Maximum Pend KWrite | Total Retries |
|------------------|-------|-----------------------|------------------------|-----------------------|------------------------|------------------|
| ----- | ----- | ----- | ----- | ----- | ----- | ----- |
| hdisk144 | 4 | 18308 | 0 | 9154 | 0 | 498 |
| hdisk158 | 3 | 9635 | 0 | 4817 | 0 | 508 |
| hdisk173 | 0 | 73980 | 0 | 36990 | 0 | 512 |
| hdisk202 | X | 0 | 0 | 0 | 0 | 0 |
| Network Summary: | | | | | | |
| 103.17.133.102 | | 22385 | 0 | 11192 | 0 | 10 |
| 103.17.133.104 | | 18614 | 0 | 9306 | 0 | 201 |
| 103.17.133.202 | | 43625 | 0 | 21811 | 0 | 1884 |
| 103.17.133.204 | | 17410 | 0 | 8704 | 0 | 32 |

Example 8

Run the **rpvstat** command with the **-m** and **-n** options to display maximum pending statistics (high water marks) with breakout by network. This displays the high water mark statistics first by RPV device (with breakout by network) and then by network (for all devices).

```
>>>rpvstat -m -n
```

| RPV Client | cx | Maximum Pend Reads | Maximum Pend Writes | Maximum Pend KRead | Maximum Pend KWrite | Total Retries |
|----------------|-------|-----------------------|------------------------|-----------------------|------------------------|------------------|
| ----- | ----- | ----- | ----- | ----- | ----- | ----- |
| hdisk144 | 4 | 18308 | 0 | 9154 | 0 | 498 |
| 103.17.133.102 | Y | 4111 | 0 | 2221 | 0 | 124 |
| 103.17.133.104 | Y | 3903 | 0 | 2843 | 0 | 101 |
| 103.17.133.202 | Y | 5384 | 0 | 2045 | 0 | 98 |
| 103.17.133.204 | Y | 4910 | 0 | 2045 | 0 | 175 |
| hdisk158 | 3 | 9635 | 0 | 4817 | 0 | 508 |
| 103.17.133.102 | Y | 2400 | 0 | 1201 | 0 | 123 |
| 103.17.133.104 | Y | 2300 | 0 | 1210 | 0 | 120 |
| 103.17.133.202 | N | 2700 | 0 | 1204 | 0 | 99 |

| | | | | | | |
|------------------|---|-------|---|-------|---|-----|
| 103.17.133.204 | Y | 2235 | 0 | 1202 | 0 | 166 |
| hdisk173 | 0 | 73980 | 0 | 36990 | 0 | 512 |
| 103.17.133.102 | N | 21101 | 0 | 9384 | 0 | 120 |
| 103.17.133.104 | N | 19872 | 0 | 9034 | 0 | 110 |
| 103.17.133.202 | N | 16004 | 0 | 9110 | 0 | 159 |
| 103.17.133.204 | N | 17003 | 0 | 9462 | 0 | 123 |
| hdisk202 | X | 0 | 0 | 0 | 0 | 0 |
| Network Summary: | | | | | | |
| 103.17.133.102 | | 27612 | 0 | 12806 | 0 | 367 |
| 103.17.133.104 | | 52150 | 0 | 13087 | 0 | 331 |
| 103.17.133.202 | | 24088 | 0 | 12359 | 0 | 356 |
| 103.17.133.204 | | 24148 | 0 | 12708 | 0 | 464 |

Example 9

Run the **rpvstat** command with **-A** option. This displays overall statistics for asynchronous mirroring plus statistics for all devices.

```
>>>rpvstat -A
```

| RPV Client | ax | Completed Async Writes | Completed Async KB Writes | Cached Async Writes | Cached Async KB Writes | Pending Async Writes | Pending Async KB Writes |
|------------|----|------------------------|---------------------------|---------------------|------------------------|----------------------|-------------------------|
| hdisk15 | A | 73210 | 92434343 | 35 | 63423 | 15 | 15478 |
| hdisk16 | A | 63465 | 74979797 | 69 | 94232 | 2 | 1236 |
| hdisk17 | A | 1095781 | 8956223105 | 52 | 76912 | 4 | 5600 |

Example 10

Run the **rpvstat** command with **-C** option. This displays the asynchronous I/O cache statistics.

```
>>>rpvstat -C
```

| GMVG Name | ax | Total Async Writes | Max Cache Util % | Pending Cache Writes | Total Cache Wait % | Max Cache Wait | Cache Free Space KB |
|-----------|----|--------------------|------------------|----------------------|--------------------|----------------|---------------------|
| glvm_vg | A | 22211737 | 89.80 | 234444 | 1.05 | 360 | 2234568 |

The gmvstat command

The **gmvstat** command provides status monitoring for Geographically Mirrored Volume Groups (GMVGs) to display the information for one or more GMVGs.

These GMVGs include:

- GMVG name
- Number of physical volumes (PVs) in the GMVG on the local system
- Number of remote physical volumes (RPVs) in the GMVG represented on the local system but physically residing on a remote system
- Total number of volumes (PVs + RPVs)
- Number of stale volumes
- Total number of physical partitions (PPs) in the volume group
- Number of stale PPs in the volume group
- Synchronization state of the GMVG - percentage of PPs synchronized (that is, not stale)

The **gmvstat** command allows this information to be displayed for all GMVGs on the system or for a subset of GMVGs specified by GMVG name on the command line. The command display optionally includes the associated **rpvstat** display output for RPV clients associated with the specified GMVG(s).

Sample display output for gmvgstat command

These samples display various output for the **gmvgstat** command.

Example 1

Run the **gmvgstat** command with no options to display each GMVG on the local node along with associated statistics.

```
>>>gmvgstat
GMVG Name      PVs  RPVs  Tot Vols  St Vols  Total PPs  Stale PPs  Sync
-----
red_gmv7       60   63    123      0        29846140   0          100%
blue_gmv23     5    5     10       1         5926      384        93%
green_gmv19   152  152   304     152       91504     45752     50%
```

Example 2

Run the **gmvgstat** command with no options, but specifying the GMVG blue_gmv23 to display statistics for that particular GMVG only.

```
>>>gmvgstat blue_gmv23
GMVG Name      PVs  RPVs  Tot Vols  St Vols  Total PPs  Stale PPs  Sync
-----
blue_gmv23     5    5     10       1         235       10         95%
```

Example 3

Run the **gmvgstat** command with the -t and -r options, specifying the GMVG blue_gmv23 to display statistics for the specified GMVG followed by statistics for each RPV included in blue_gmv23 (from the **rpvstat** command).

```
>>>gmvgstat -t -r blue_gmv23
Geographically Mirrored Volume Group Information      02:23:57 PM 16 Feb 2007
-----
c689n02.ppd.pok.ibm.com
siteA
GMVG Name      PVs  RPVs  Tot Vols  St Vols  Total PPs  Stale PPs  Sync
-----
blue_gmv23     3    3     6         1         235       10         95%

Remote Physical Volume Statistics:

RPV Client      Comp Reads  Comp Writes  Comp KPread  Comp KBWrite  Errors
-----
cx  Pend Reads  Pend Writes  Pend KRead   Pend KBWrite
-----
hdisk144        2           543         0             5430         0           0
                0           0           0             0            0
hdisk158        2           228         0             2280         0           0
                0           0           0             0            0
hdisk177        2           428         0             4280         0           2
                1           0           0             10           0
```

Example 4

Run the **gmvgstat** command to display statistics for the two GMVGs red_gmv7 and green_gmv19. Monitor these GMVGs every 60 seconds for a total of 4 times and show the date and time with each display.

```
>>>gmvgstat -t -i 60 -c 4 red_gmv7 green_gmv19
Geographically Mirrored Volume Group Information      02:23:57 PM 16 Feb 2007
-----
c689n02.ppd.pok.ibm.com
siteA
GMVG Name      PVs  RPVs  Tot Vols  St Vols  Total PPs  Stale PPs  Sync
-----
red_gmv7       60   63    123      0        29846140   0          100%
green_gmv19   152  152   304     152       91504     45752     50%

Geographically Mirrored Volume Group Information      02:24:58 PM 16 Feb 2007
-----
c689n02.ppd.pok.ibm.com
siteA
GMVG Name      PVs  RPVs  Tot Vols  St Vols  Total PPs  Stale PPs  Sync
-----
```

| | | | | | | | |
|--------------------------------------------------|-------|-------|----------|---------|-------------------------------------------------------------|-----------|-------|
| red_gmvg7 | 60 | 63 | 123 | 0 | 29846140 | 0 | 100% |
| green_gmvg19 | 152 | 152 | 304 | 99 | 91504 | 29799 | 67% |
| Geographically Mirrored Volume Group Information | | | | | 02:25:59 PM 16 Feb 2007 c689n02.ppd.pok.ibm.com siteA | | |
| ----- | ----- | ----- | ----- | ----- | ----- | ----- | ----- |
| GMVG Name | PVs | RPVs | Tot Vols | St Vols | Total PPs | Stale PPs | Sync |
| red_gmvg7 | 60 | 63 | 123 | 0 | 29846140 | 0 | 100% |
| green_gmvg19 | 152 | 152 | 304 | 10 | 91504 | 3010 | 96% |
| Geographically Mirrored Volume Group Information | | | | | 02:27:00 PM 16 Feb 2007 c689n02.ppd.pok.ibm.com siteA | | |
| ----- | ----- | ----- | ----- | ----- | ----- | ----- | ----- |
| GMVG Name | PVs | RPVs | Tot Vols | St Vols | Total PPs | Stale PPs | Sync |
| red_gmvg7 | 60 | 63 | 123 | 0 | 29846140 | 0 | 100% |
| green_gmvg19 | 152 | 152 | 304 | 0 | 91504 | 0 | 100% |

SMIT interfaces

The **rpvstat** and **gmvgstat** commands can also be run from SMIT.

The descriptions of their SMIT interfaces are provided in the following sections.

SMIT interfaces for GLVM status monitoring tools - rpvstat and gmvgstat

The main entry point for the GLVM Status Monitor commands is located on the Geographical Logical Volume Manager Utilities panel (fastpath: `glvm_utils`). The SMIT menu item is Status Monitors:

```

Geographical Logical Volume Manager Utilities

Move cursor to desired item and press Enter.

Geographically Mirrored Volume Groups
Geographically Mirrored Logical Volumes
Remote Physical Volume Clients
Remote Physical Volume Server
Status Monitors

F1=Help          F2=Refresh      F3=Cancel      F8=Image
F9=Shell         F10=Exit       Enter=Do

```

Selecting Status Monitors brings up the next panel (fastpath: `glvmmonitors`):

```

Status Monitors

Move cursor to desired item and press Enter.

Remote Physical Volume Client Statistics
Geographically Mirrored Volume Group Status
Asynchronous Mirroring Statistics

F1=Help          F2=Refresh      F3=Cancel      F8=Image
F9=Shell         F10=Exit       Enter=Do

```

Select Remote Physical Volume Client Status to display the SMIT interface for the **rpvstat** command for physical volumes. Select Geographically Mirrored Volume Group Status to display the SMIT interface for the **gmvgstat** command. Select Asynchronous Mirroring Statistics to display the SMIT interface for the **rpvstat** command for asynchronous monitoring.

SMIT interface for rpvstat

These sample SMIT interfaces describe using the **rpvstat** command.

Select Remote Physical Volume Client Statistics from the Status Monitors panel to display the main SMIT interface panel for the **rpvstat** command (fastpath: rpvstat).

```

Remote Physical Volume Client Statistics

Move cursor to desired item and press Enter.

Display Remote Physical Volume Client Statistics
Display Remote Physical Volume Client Statistics by Network
Display Remote Physical Volume Client Maximum Pending Statistics
Reset Remote Physical Volume Client Statistics

F1=Help          F2=Refresh       F3=Cancel        F8=Image
F9=Shell         F10=Exit         Enter=Do

```

Select Display Remote Physical Volume Client Statistics to display the following panel (fastpath: rpvstat_dialog):

```

Display Remote Physical Volume Client Statistics

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

Specific RPV(s) to display (leave blank for all)      [Entry Fields]
Display statistics for individual mirroring networks?  no          +
Include date and time in display?                    no          +
Monitor interval in seconds                          []          #
Number of times to monitor                           []          #
Display applicable monitored statistics as delta values? no          +

F1=Help          F2=Refresh       F3=Cancel        F4=List
F5=Reset         F6=Command       F7=Edit          F8=Image
F5=Shell         F10=Exit         Enter=Do

```

| <i>Display Remote Physical Volume Client Statistics fields</i> | |
|----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Specific RPV(s) to display (leave blank for all) | You can leave this field empty to display information for all RPV clients. Pressing PF4 provides a list from which you can use PF7 to select one or more RPV clients for which to display information. Or, you can manually enter the names of one or more RPV clients in the entry field. |
| Display statistics for individual mirroring networks? | You can alternately select no or yes. |
| Include date and time in display? | You can alternately select no or yes. |
| Monitor interval in seconds | This field is optional and requires an integer value greater than zero and less than or equal to 3600. |
| Number of times to monitor | This field is optional, but requires that the Monitor interval in seconds field have a value. The Number of times to monitor field requires an integer value greater than 0 and less than or equal to 999999. |
| Display applicable monitored statistics as delta values? | You can alternately select no or yes. This field applies only if the Monitor interval in seconds field has a value. |

After the fields are completed, press Enter to run the **rpvstat** command to display statistical information for all of the indicated RPV clients.

Select Display Remote Physical Volume Client Statistics by Network on the previous Remote Physical Volume Client Statistics panel to display the following panel (fastpath: rpvstat_net_dialog):

```

Display Remote Physical Volume Client Statistics by Network

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
Include date and time in display?      no      +
Monitor interval in seconds            []      #
Number of times to monitor              []      #
Display applicable monitored statistics as delta values?  no      +

F1=Help      F2=Refresh      F3=Cancel      F4=List
F5=Reset     F6=Command     F7=Edit       F8=Image
F5=Shell     F10=Exit       Enter=Do

```

| <i>Display Remote Physical Volume Client Statistics by Network fields</i> | |
|---------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Include date and time in display? | You can alternately select no or yes. |
| Monitor interval in seconds | This field is optional. This field requires an integer value greater than zero and less than or equal to 3600. |
| Number of times to monitor | This field is optional, but requires the Include date and time in display field to have a value. The Monitor interval in seconds field requires an integer value greater than 0 and less than or equal to 999999. |
| Display applicable monitored statistics as delta values? | You can alternately select no or yes. This field applies only if the Monitor interval in seconds field has a value. |

After the fields are completed, press Enter to run the **rpvstat -N** command to display statistical information for all of the remote mirroring networks.

Select Display Remote Physical Volume Client Maximum Pending Statistics on the previous Remote Physical Volume Client Statistics panel to display the following panel (fastpath: rpvstat_pending_dialog):

```

Display Remote Physical Volume Client Maximum Pending Statistics

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
Specific RPV(s) to display (leave blank for all)      []      +
Display statistics for individual mirroring networks  no      +
Include date and time in display?                      no      +

F1=Help      F2=Refresh      F3=Cancel      F4=List
F5=Reset     F6=Command     F7=Edit       F8=Image
F5=Shell     F10=Exit       Enter=Do

```

| <i>Display Remote Physical Volume Client Maximum Pending Statistics fields</i> | |
|--------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Specific RPV(s) to display (leave blank for all) | You can leave this field blank to display statistics for all RPV clients. Pressing PF4 provides a list from which you can use PF7 to select one or |

| Field | Value |
|------------------------------------------------------|---------------------------------------------------------------------------------------------------|
| | more RPV clients. You can also manually enter the names of one or more RPV clients in this field. |
| Display statistics for individual mirroring networks | You can alternately select no or yes. |
| Include date and time in display? | You can alternately select no or yes. |

After the fields are completed, press Enter to run the **rpvstatus -m** command to display high water mark statistical information for pending statistics.

Select Reset RPV Client statistics on the previous Remote Physical Volume Client Statistics panel to display the following panel (fastpath: rpvstat_reset_dialog):

```

Reset Remote Physical Volume Client Statistics

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]

Specific RPV(s) to reset (leave blank for all)  []          +

F1=Help          F2=Refresh          F3=Cancel          F4=List
F5=Reset         F6=Command         F7=Edit            F8=Image
F5=Shell         F10=Exit           Enter=Do

```

| <i>Reset Remote Physical Volume Client Statistics</i> | |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Specific RPV(s) to reset (leave blank for all) | You can leave this field blank to reset statistics for all RPV clients. Pressing PF4 provides a list from which you can use PF7 to select one or more RPV clients. You can also manually enter the names of one or more RPV clients in this field. |

Press Enter to run the **rpvstat -R** command to reset the statistical counters in the indicated RPV clients.

SMIT interface for rpvstat for asynchronous mirroring

These sample SMIT interfaces describe using the **rpvstat** command for asynchronous mirroring.

If you select **Asynchronous Mirroring Statistics**, the following panel is displayed:

```

Asynchronous Mirroring Statistics

Move cursor to desired item and press Enter.

Display Asynchronous IO Progress Statistics
Display Asynchronous IO Cache Statistics
Reset Asynchronous IO Statistics

F1=Help          F2=Refresh          F3=Cancel          F8=Image
F9=Shell         F10=Exit           Enter=Do

```

If you select **Display Asynchronous IO Progress Statistics**, the following panel is displayed:

```

Display Asynchronous IO Progress Statistics

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

Specific RPV(s) to display (leave blank for all)  [Entry Fields]
Include date and time in display?                []          +
Monitor interval in seconds                       no          +
                                                    []          #

```

```

Number of times to monitor          []          #
Display applicable monitored statistics as delta values?      no          +

F1=Help          F2=Refresh          F3=Cancel          F8=Image
F9=Shell         F10=Exit           Enter=Do

```

| <i>Display Asynchronous IO Progress Statistics</i> | |
|----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Specific RPV(s) to display (leave blank for all) | You can leave this field empty to display information for all RPV clients. Pressing PF4 provides a list from which you can use PF7 to select one or more RPV clients for which to display information. Or, you can manually enter the names of one or more RPV clients in the entry field. |
| Include date and time in display? | You can alternately select no or yes. |
| Monitor interval in seconds | This field is optional and requires an integer value greater than zero and less than or equal to 3600. |
| Number of times to monitor | This field is optional, but requires that the Monitor interval in seconds field have a value. The Number of times to monitor field requires an integer value greater than 0 and less than or equal to 999999. |
| Display applicable monitored statistics as delta values? | You can alternately select no or yes. This field applies only if the Monitor interval in seconds field has a value. |

If you select **Display Asynchronous IO Cache Statistics**, the following panel is displayed:

```

Display Asynchronous IO Cache Statistics

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

Include date and time in display?          [Entry Fields]
Monitor interval in seconds                no          +
Number of times to monitor                 [2]         #
Display applicable monitored statistics as delta values?  [5]         #
                                           yes         +

F1=Help          F2=Refresh          F3=Cancel          F8=Image
F9=Shell         F10=Exit           Enter=Do

```

| <i>Display Asynchronous IO Cache Statistics fields</i> | |
|----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Include date and time in display? | You can alternately select no or yes. |
| Monitor interval in seconds | This field is optional and requires an integer value greater than zero and less than or equal to 3600. |
| Number of times to monitor | This field is optional, but requires that the Monitor interval in seconds field have a value. The Number of times to monitor field requires an integer value greater than 0 and less than or equal to 999999. |
| Display applicable monitored statistics as delta values? | You can alternately select no or yes. This field applies only if the Monitor interval in seconds field has a value. |

Selecting **Reset Asynchronous IO Statistics** runs the `rpvstat -A -R` and `rpvstat -C -R` commands to reset all statistics.

SMIT interface for gmvgstat

These sample SMIT interfaces describe using the **gmvgstat** command.

Select Geographically Mirrored Volume Group Status from the new Status Monitors panel to display the next panel (fastpath: `gmvgstat`).

Geographically Mirrored Volume Group Status

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

| | | [Entry Fields] | |
|---------------------------------------------------|--|----------------|---|
| Specific GMVG(s) to display (leave blank for all) | | [] | + |
| Include associated RPV Client statistics? | | no | + |
| Include header with display? | | no | + |
| Monitor interval in seconds | | [] | # |
| Number of times to monitor | | [] | # |

| | | | |
|----------|------------|-----------|----------|
| F1=Help | F2=Refresh | F3=Cancel | F4=List |
| F5=Reset | F6=Command | F7=Edit | F8=Image |
| F5=Shell | F10=Exit | Enter=Do | |

| <i>Geographically Mirrored Volume Group Status fields</i> | |
|-----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Field | Value |
| Specific RPV(s) to display (leave blank for all) | You can leave this field empty to display information for all valid, available, online GMVGs. Pressing PF4 provides a list from which you can use PF7 to select one or more GMVGs for which to display information. You can also manually enter the names of one or more GMVGs in this field. |
| Include associated RPV Client statistics? | You can alternately select no or yes. |
| Include header with display? | You can alternately select no or yes. |
| Monitor interval in seconds | This field is optional and requires an integer value greater than zero and less than or equal to 3600. |
| Number of times to monitor | This field is optional, but requires a value in the Monitor interval in seconds field. The Number of times to monitor field requires an integer value greater than 0 and less than or equal to 999999. |

After the fields are complete, press Enter to run the **gmvgsstat** command to display statistical information for all of the indicated GMVGs.

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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM® under terms of the IBM® Customer Agreement, IBM® International Program License Agreement or any equivalent agreement between us.

The performance data 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 [Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml) at www.ibm.com/legal/copytrade.shtml.

Index

A

activating

- asynchronous mirroring [64](#)
- remote physical volumes [73](#)
- remote physical volumes
 - defined clients [75](#)
- volume group
 - after production site failure [65](#)
 - after production site recovery [65](#)

adding

- remote physical volumes
 - to mirror pool [60](#)
 - to volume group [57](#)
- remote site mirror copy
 - to logical volume [66](#)
 - to mirror pool [61](#)
 - to volume group [57](#)

asynchronous mirroring

- activating [64](#)
- best practices [34](#)
- changing attributes [63](#)
- configuration overview [47](#)
- configuring [50](#)
- converting to [62](#)
- listing for a mirror pool [62](#)
- overview [14](#)
- planning [32](#)
- recovering from disaster [101](#)

B

best practices

- asynchronous mirroring [34](#)

C

cache device failure [93](#)

CD-ROM [38](#)

changing

- asynchronous mirroring attributes
 - for a mirror pool [63](#)
- remote physical volumes
 - clients [73](#)
 - multiple clients [74](#)
 - multiple servers [71](#)
 - server [71](#)
 - site name [70](#)

cluster

- integrating GLVM [80](#)

concepts

- data mirroring [8](#)
- Geographic Logical Volume Manager [7](#)
- GLVM [10](#)
- remote physical volume [11](#)
- stand-alone GLVM [12](#)
- synchronous mirroring [13](#)

configuring

- asynchronous mirroring [50](#)
- GLVM [41](#)
- logical volumes [56](#)
- mirror copies [56](#)
- prerequisites [41](#)
- remote physical volumes [52, 56](#)
- resource groups [83](#)
- sites [81](#)
- verifying GLVM configuration [89](#)
- XD-type network [81](#)

converting

- asynchronous mirroring
 - for a mirror pool [62](#)
- synchronous mirroring
 - for a mirror pool [63](#)

D

data divergence [14, 101](#)

data divergence

- overriding default [102](#)

data mirroring

- concepts [8](#)

F

fallback

- GLVM [16](#)

fallover [100](#)

fallover

- GLVM [16](#)

G

Geographic Logical Volume Manager

- concepts [7](#)

GLVM [104](#)

GLVM

- asynchronous mirroring [14](#)
- cache device failure [93](#)
- concepts [7, 10](#)
- configuration prerequisites [41](#)
- configuring [41](#)
- fallback [16](#)
- fallover [16](#)
- forced varyon [95](#)
- installation scenarios [36](#)
- installing [35](#)
- installing
 - from CD-ROM [38](#)
 - from hard disk [37](#)
 - from installation server [37](#)
 - troubleshooting [40](#)
 - verification [39](#)
- installing prerequisites [36](#)

- integrating 80
- limitations 24
- maintaining 91, 93
- monitoring 104
- quorum 95
- requirements 24
- sample configuration 19
- selective fallover 100
- sites 16
- synchronous mirroring 13
- troubleshooting 93
- upgrading 40
- verifying configuration 89

GLVM Configuration Assistant 42

GLVM Configuration Assistant

- asynchronous 43
- synchronous 44
- Troubleshoot 93

gmvgsat command 109

gmvgsat command

- SMIT interfaces 111, 115

H

hard disk 37

I

installation server 37

installing

- from CD-ROM 38
- from hard disk 37
- from installation server 37
- GLVM 35
- prerequisites 36
- scenarios 36
- troubleshooting 40
- verifying 39

integrating

GLVM

- cluster 80

L

listing

- asynchronous mirroring
 - for a mirror pool 62
- logical volumes 69
- remote physical volumes
 - clients 73
- volume groups 77

logging messages 93

logical volume

- adding remote site mirror copy 66
- configuring 56
- listing 69
- verifying mirror copy location 69

M

maintaining

- GLVM 91, 93

managing

- mirror copy 56
- mirror pools 60

mirror copy

- adding to logical volume 66
- adding to volume group 57
- configuring 56
- managing 56
- removing from a logical volume 77
- removing from mirror pool 61
- removing from volume group 59
- verifying for a volume group 59
- verifying locations 69

mirror pool

- activating asynchronous mirroring 64
- adding remote physical volume 60
- adding remote site mirror copy 61
- changing asynchronous mirroring attributes 63
- converting to asynchronous mirroring 62
- converting to synchronous mirroring 63
- listing asynchronous mirroring 62
- managing 60
- mirror copy**
 - adding to mirror pool 61
 - planning 33
 - removing remote physical volume 60
 - removing remote site mirror copy 61

monitoring 104

N

network

- configuring XD-type 81
- planning 25
- planning**
 - security 26

O

overview

- asynchronous mirror
 - configuration 47
- asynchronous mirroring 14

P

planning

- asynchronous mirroring 32
- limitations 24
- mirror pools 33
- network security 26
- networks 25
- remote physical volumes 25
- requirements 24
- resource groups 27
- sites 25
- volume groups 25

Q

quorum 95

R

recovering 101

remote physical volume

activating 73

activating defined clients 75

adding to mirror pool 60

adding to volume group 57

changing clients 73

changing multiple clients 74

changing multiple servers 71

changing server 71

changing site name 70

concepts 11

configuring 52, 56

listing clients 73

planning 25

removing 72

removing clients 75

removing from a volume group 76

removing from mirror pool 60

replacing 76

showing server 71

remote site mirror copy

adding to logical volume 66

adding to mirror pool 61

adding to volume group 57

removing from logical volume 77

removing from mirror pool 61

removing from volume group 59

verifying for a volume group 59

removing

remote physical volumes

clients 75

from a volume group 76

from mirror pool 60

servers 72

remote site mirror copy

from a logical volume 77

from mirror pool 61

from volume group 59

replacing

remote physical volumes 76

resource group

configuring 83

planning 27

rpvstat command 104, 114

rpvstat command

SMIT interfaces 112, 111

SMIT interfaces

asynchronous mirroring 114

S

sample configuration

GLVM 19

showing

remote physical volume

clients 73

remote physical volumes

clients 73

server 71

site

changing 70

configuring 81

GLVM 16

planning 25

remote physical volume site name 70

SMIT interfaces

asynchronous mirroring 114

gmvstat command 111, 115

rpvstat command 112, 111

rpvstat command

asynchronous mirroring 114

stand-alone

concepts 12

synchronizing

GLVM configuration 89

synchronous mirroring

concepts 13

converting 63

T

troubleshooting

GLVM 93

installation 40

logging messages 93

U

upgrading

GLVM 40

V

verifying

GLVM configuration 89

installation 39

mirror copy

for a volume group 59

mirror copy locations

logical volume 69

volume group

activating after production site failure 65

activating after production site recovery 65

listing 77

planning 25

removing remote physical volumes 76

© Copyright International Business Machines Corporation 2017, 2022

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

