HP StorageWorks

P4000 Multi-Site HA/DR Solution Pack user guide

Abstract

This guide contains detailed instructions for designing and implementing the Multi-Site SAN features of the HP P4000 SAN Solution. The Multi-Site SAN features enable you to synchronously and automatically mirror data between geographic sites. Topics include designs for network configurations and instructions about implementing the network designs. The intended audience is system administrators who manage HP P4000 SAN solutions.

Part number: AX696-96004
Third edition: March 2010

Legal and notice information

Acknowledgements

Microsoft, Windows, Windows XP, and Windows NT are U.S. registered trademarks of Microsoft Corporation.

Contents

Figures

Tables

1 Designing a Multi-Site SAN

The Multi-Site SAN features enable you to synchronously mirror data between sites automatically. You create a Multi-Site SAN by configuring equal numbers of storage nodes into sites in the software. The Multi-Site SAN software enables a variety of features that are based on geographic awareness and synchronous replication:

- Multi-Site clusters that can:
  - Span up to three sites
  - Span multiple subnets
  - Have multiple VIPs
- Geographical awareness, enabled by designating storage nodes as members of a particular site.
- Synchronously replicated (mirrored) data between sites, based on volume data protection level.
- Site information that ensures that data is mirrored between sites for high availability and disaster recovery.
- I/O path preferencing, so that application servers connect to storage nodes located in the same site.
- Failover Manager support for automatic failover/failback in three-site configurations without requiring a physical storage node in the third site.
- Failover Manager for quorum management if the network connection between the two sites becomes unavailable.
- Recover quorum via CLI when a site failure also results in losing quorum in the management group.

Failover Manager overview

The Failover Manager is a specialized version of the SAN/iQ software, designed to run as a virtual appliance in a VMware environment. The Failover Manager participates in the management group as a real manager in the system; however, it performs quorum operations only, not data movement operations. It is especially useful in a Multi-Site SAN configuration to manage quorum for the multi-site configuration without requiring additional physical hardware at a site.

The Failover Manager is supported on VMware Server, Workstation, and Player. For detailed information about installing and configuring the Failover Manager, see "Using Specialized Managers" in the *HP StorageWorks SAN Solutions Storage System User Guide*.

Requirements

- A Multi-Site SAN requires a feature key for each storage node in the management group that is also participating in a Multi-Site cluster.
- All sites must have the same number of storage nodes. Storage nodes must be added to or removed from Multi-Site clusters equally.

Designing a Multi-Site SAN

The Multi-Site SAN software offers multiple features and the design flexibility to protect against certain types of failures in the environment.

Data replication, site design, and quorum managers combine to protect against certain types of failures in the environment. Some of the common types of failures a Multi-Site SAN protects against include data center failures and storage node failures:

Data center failures

- Site power outage
- Site network outage
- Site disaster (fire, flood, terrorist attack, and so forth)

Individual storage node failures

- Power outage
- Network outage
- Node disaster
- System failure (backplane, motherboard, RAM, CPU)

Configuring data replication for optimal availability and redundancy

The SAN/iQ software offers six levels of synchronous replication at a volume level. In order to protect the data across sites, you must choose a data protection level with a number of mirrors that is at least equal to the number of sites in the configuration. For instance, if the storage cluster spans three sites, then volumes must be configured as Network RAID–10+1 in order to insure that data is the same on each site.

NOTE:

Network RAID-0, Network RAID-5, and Network RAID-6 will not tolerate site failures.

Table 1 on page 8 describes the data replication levels for the Multi-Site SAN and associated supported configurations.

Table 1 Data protection levels and the supported number of sites

Data protection level	Number of sites supported
Network RAID-10	2
Network RAID–10+1	3 NOTE: 2 sites are supported, but not considered a best practice configuration
Network RAID–10+2	2 NOTE: 3 sites are supported, but are not considered a best practice configuration

When to use Network RAID-10+2

Using Network RAID-10+2 offers a level of data protection that enables the Multi-Site SAN to sustain a complete failure at one site and still maintain replicated data at the other site. For instance, with Network RAID-10+2 there are four copies of data for the volume on the SAN. In a two-site configuration, the Multi-Site SAN software puts two of the four copies in one site and the remaining two in the other. If a site fails, the remaining site still has two copies and remains fault-tolerant.

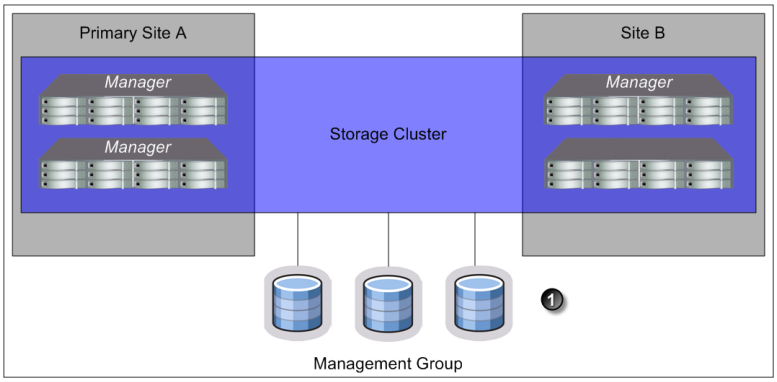

1. Volumes are configured with Network RAID-10+2

Figure 1 Sample configuration for two-site, data protection level Network RAID-10+2

Configuring sites and quorum managers to achieve optimal high availability and redundancy

Several recommended site configurations can be used with the Multi-Site SAN software. These configurations provide varying levels of availability and redundancy. Table 2 on page 10 shows the recommended configurations. Select the one that best matches your environment.

In addition to setting up the Multi-Site clusters for data replication, it is important to set up the SAN managers for quorum correctly to ensure that in the event of a failure the system can be properly recovered. Maintaining quorum can be either an automated process or manually recovered by an Administrator. Table 2 summarizes common Multi-Site SAN and manager configurations.

Table 2 Common configurations of Multi-Site SANs and mangers

Number of sites	Cluster/site setup	Total number of managers of all types	Failover Manager	Site types
three physical sites	Cluster spans two sites	Min = three; Max = five	In the third site	All sites are peer.
two physical sites, with logical third site	Cluster spans two sites	Min = three; Max = five	In the logical third site	All sites are peer.
three physical sites	Cluster spans three sites	Min = three; Max = six	Not used	All sites are peer.
two physical sites	Cluster spans two sites	Min = three; Max = five	Not used	One site is designated primary.

Best practice

In a cluster that spans multiple sites, run the same number of managers in each site. In a cluster that spans two sites, run a Failover Manager in a third site (physical or logical) that is not included in the cluster (Figure 2 on page 11). The one exception to this rule is the two-site configuration in which one site is configured as primary (Figure 5 on page 12). In this case, run a majority of the managers in the primary site.

Sample configurations

Figure 2 on page 11, Figure 3 on page 11, Figure 4 on page 12, and Figure 5 on page 12 illustrate best practice for four configurations. Note that management groups and Multi-Site clusters can have many more storage nodes than are depicted here. These graphics however, illustrate the significant points of site-to-cluster mappings and the appropriate placement of managers.

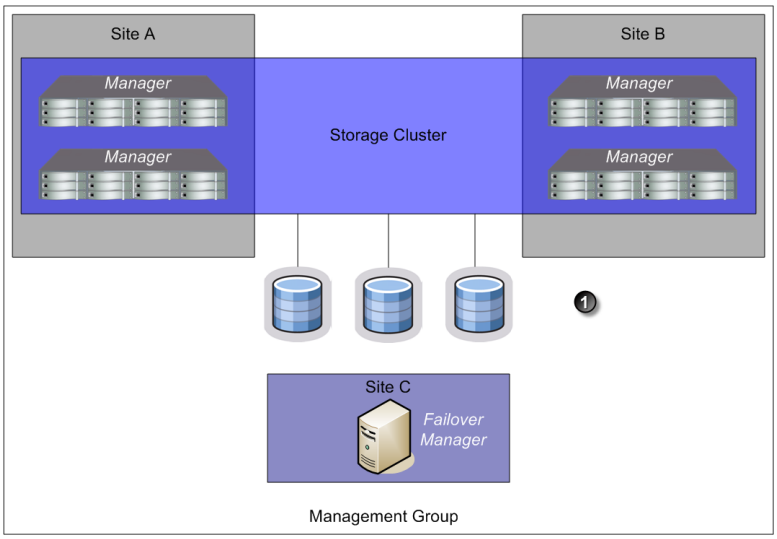

1. Volumes are configured with Network RAID-10

Figure 2 Three physical sites: cluster spans two sites, Failover Manager in third site

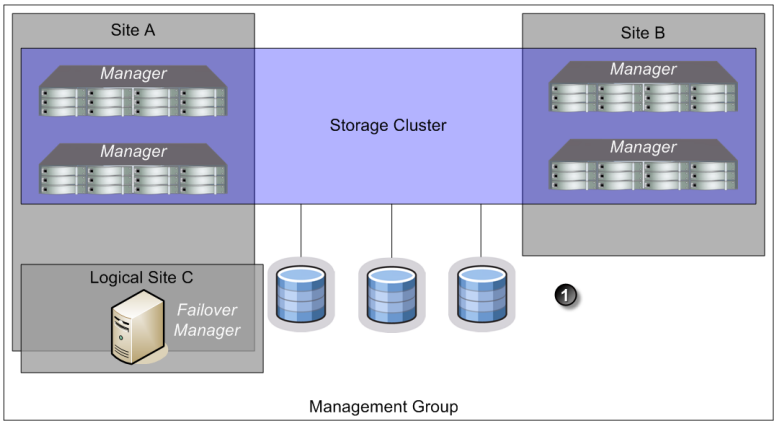

1. Volumes are configured with Network RAID-10

Figure 3 Two physical sites, logical third site: cluster spans two sites, Failover Manager in logical third site

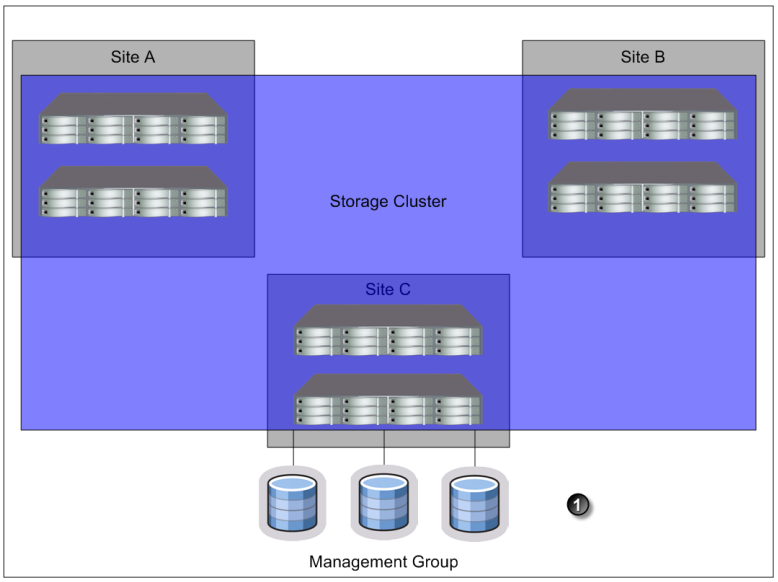

1. Volumes are configured with Network RAID-10+1

Figure 4 Three physical sites: cluster spans three sites

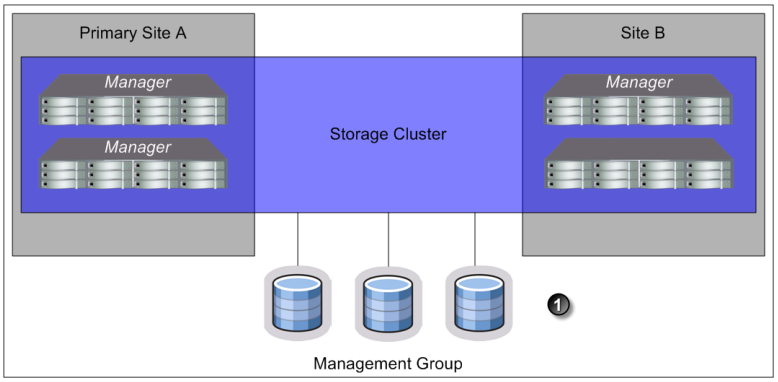

1. Volumes are configured with Network RAID-10

Figure 5 Two physical sites: cluster spans two sites, one site is the primary site

Alternative setups

The Multi-Site SAN software supports many other valid configurations than the four illustrated in the prior section.Table 3 on page 13 lists other possible configurations.

Table 3 Alternative Multi-Site SAN configurations

Alternate configuration	Description	Details
More than one Multi-Site cluster in the management group	Multiple clusters running within the same management group.	The Multi-Site SAN software supports running many clusters in the same management group. You simply need to make sure that clusters have the same number of storage nodes in each site for that particular cluster.
More than three sites	There are more than three sites in the cluster.	A particular cluster can only span up to three sites. However, the Multi-Site SAN software supports having more than one cluster. The different clusters can span different sites within the same cluster.
Using Virtual Manager	The Virtual Manager capability remains in the product (a manual process for recovering quorum).	The Virtual Manager is used in two-site configurations, and is not as commonly used as the Failover Manager. For more information about the Virtual Manager, see the chapter "Special Managers" in the *HP StorageWorks SAN Solutions Storage System User Guide*.
Using Remote Copy with Multi-Site SAN	The cluster contains Remote Copy, which like all other SAN/iQ features, is compatible with Multi-Site SAN.	You can use Remote Copy to replicate data from a Multi-Site SAN to a Disaster Recovery or Backup site. See the chapter "Understanding and Planning Remote Copy" in the *Remote Copy User Manual*.
Multiples subnet clusters in the same site	The Multi-Site SAN features can be used to set up clusters that span multiple subnets in the same site.	The Multi-Site SAN software can be used to create clusters that span multiple networks (subnets) in larger Enterprise data centers. The storage nodes all reside in the same site. However, the cluster must be Multi-Site if it is to contain storage nodes from different subnets.

Designing the network for the Multi-Site SAN

Good network design is a critical part of setting up a Multi-Site SAN that meets your needs. A good Multi-Site SAN network encompasses the following attributes:

- Reliability
- High availability
- High performance

This section contains information about the minimum recommended network design considerations when planning a network to support a Multi-Site SAN.

Best practices for implementing the Multi-Site SAN network

- Adequate bandwidth—Plan for 50MB/sec of bandwidth for each storage node in each site. For instance, if each site contains five storage nodes, then you need 250 MB/sec throughput. In this case, that translates into two Gigabit Ethernet links, or more.
- Low round-trip latency—In order to not impact disk I/O to the application server, the round-trip latency between the sites must be no more than 2ms.
- Redundant links—Have multiple physical connections (media) between the sites for redundancy. The network should be configured so that a link failure does not cause the Multi-Site SAN to go offline.
- Use Multiple Subnets—Each site should contain its own subnet. When the Multi-Site cluster spans different subnets, other features such as MPIO can be used.
  - Optimal I/O Pathing—Application servers will perform I/O only against storage nodes that are in the same site (subnet) as the server.
  - Multiple VIPs—A VIP is assigned to each site/subnet pair. A cluster can have more than one VIP.
- Use the Failover Manager—Unless you have storage nodes in three physical sites, use a Failover Manager. Put the Failover Manager in the third site (physical or logical) located on a network that is isolated from the other sites and has a different power connection.

Common Multi-Site SAN network designs

Based on the Multi-Site SAN configuration, there are many network designs that can be implemented that solve the redundancy, high availability and performance requirements for the Multi-Site SAN. A few commonly used designs are:

- Dual redundant links between the sites
- Full-mesh triangular (three) redundant links between sites
- Hub and spoke central network core with redundant links out to each site
- Combination full-mesh core network with triangular (three) links between sites

Balancing the requirements for high availability during certain types of failures in the Multi-Site SAN and the budget associated with the network infrastructure typically dictates the network architecture a customer will choose to implement. Obviously, the most redundant network typically will cost the most to implement. However, not implementing redundancy in the Multi-Site SAN network significantly curtails the benefits of the Multi-Site SAN software. Table 4 on page 14 shows the common Multi-Site network configurations and the types of failures against which they provide protection.

Table 4 Common Multi-Site network configurations and failure protection of each

Network topology	Site failure protection	Network failure protection	Storage node failure protection
Dual redundant links between sites	Yes, with manual intervention	Yes	Yes
Full-mesh triangular links between sites	Yes	Yes	Yes
Hub & spoke central core	Yes	Yes, so long as the core does not fail	Yes
Combo full-mesh core with triangular links between sites	Yes	Yes	Yes

Sample recommended network configurations for Multi-Site SANs

Samples of common network configurations for the topologies described in Table 4 on page 14 are illustrated in this section. The clusters and sites can have many more storage nodes than are depicted here. However, the pictures are more focused on the network design.

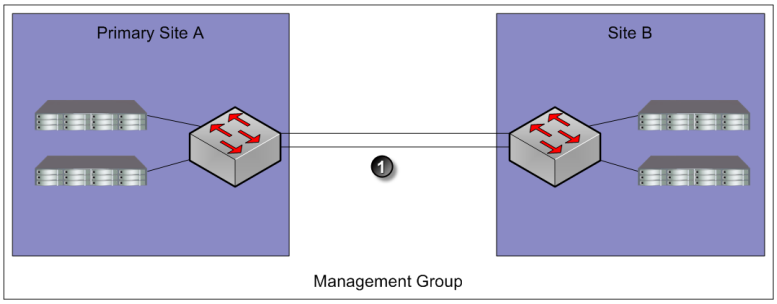

1. Dual redundant GigE links

Figure 6 Dual redundant links between two sites

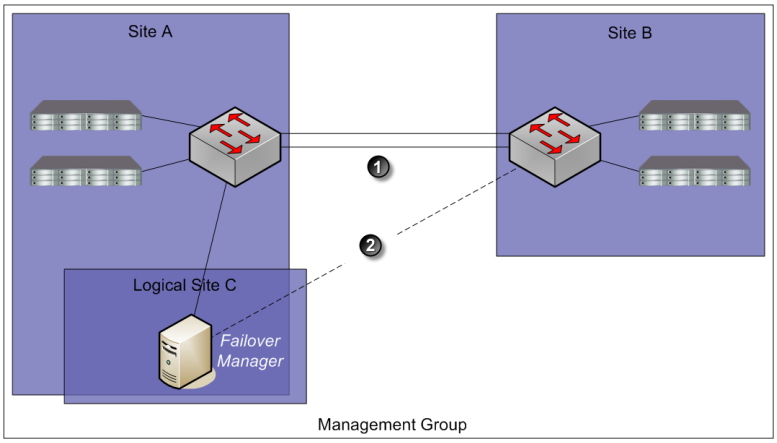

1. Dual redundant GigE links
2. Optional redundant path

Figure 7 Dual redundant links between two sites with Failover Manager in logical third site

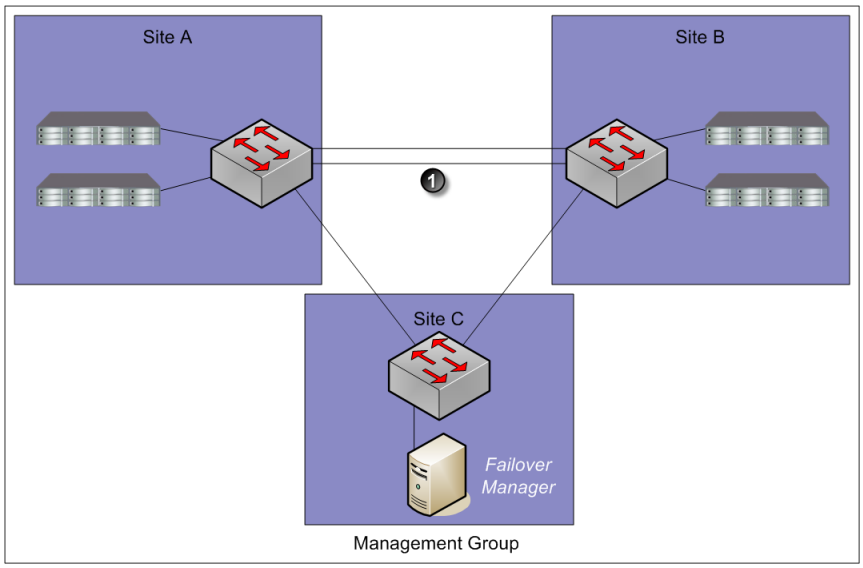

1. Dual redundant GigE links

Figure 8 Triangular network with Failover Manager

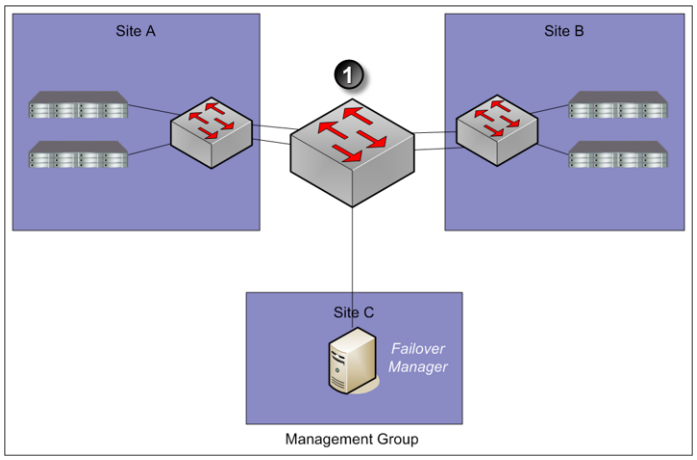

1. Core network hub and spoke

Figure 9 Hub and spoke network core with redundant links out to each site

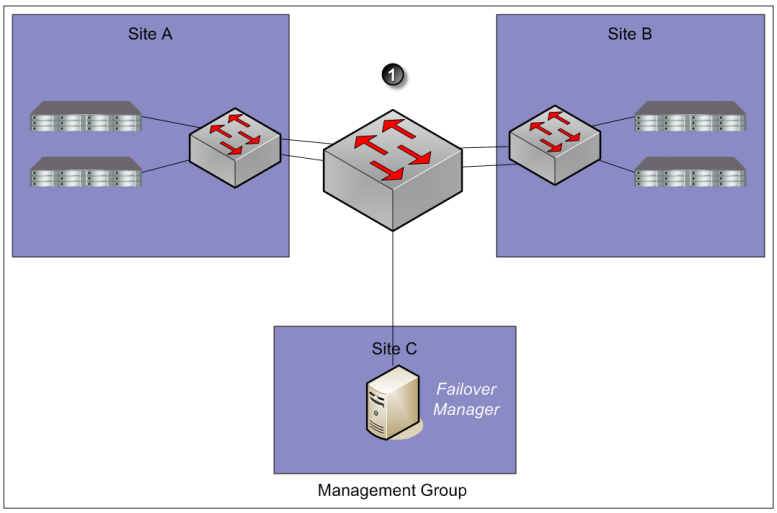

1. Core network hub & spoke

Figure 10 Combination full-mesh core with triangular links to all sites

Setting up multiple subnets for the Multi-Site SAN

Many of the advanced features in the Multi-Site SAN software depend on the implementation of multiple subnets within the iSCSI cluster. Preferably, customers should implement one subnet per site. By doing so, you can take advantage of the following Multi-Site SAN features:

- One Virtual IP address per site. iSCSI discovery sessions are terminated on a storage node that is local to the application server in that site.
- Virtual IP addresses are hosted only on storage nodes that reside in the same subnet as the VIP.
- Read and Write I/O requests are forwarded from the application server to a storage node that is in the same site as the server. This ensures that the I/O path is optimal and multiple network hops across the network link between the sites are eliminated.

Figure 11 on page 18 shows an example of how the Multi-Site SAN setup maps to subnets, VIPs, and application servers.

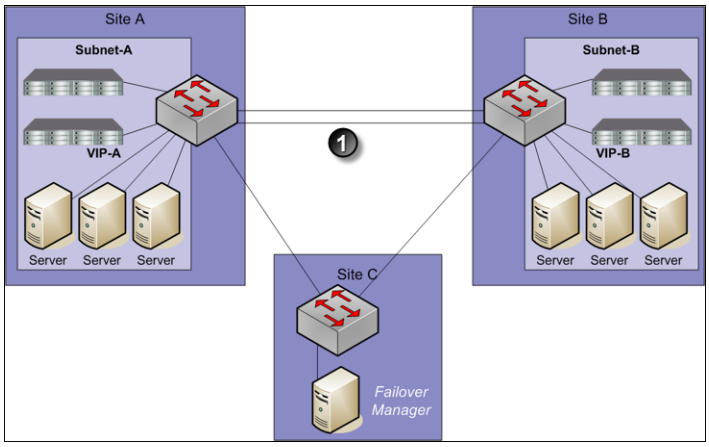

1. Dual redundant GigE links

Figure 11 Multi-Site SAN mapping to subnets, VIPs, and application servers

Using the HP LeftHand P4000 DSM for MPIO with multiple subnets

When you are using multiple subnets in a Multi-Site SAN, the HP LeftHand P4000 DSM for MPIO will only build I/O paths to the storage nodes that are in the same subnet as the application server. See the *HP StorageWorks P4000 Windows Solution Pack User Manual* for more information about HP LeftHand P4000 DSM for MPIO .

Setting up application clusters in a Multi-Site SAN environment

Just as you can set up a Multi-Site SAN for data availability even when a complete site is offline, you can set up application clusters (Microsoft Cluster Server, VMware ESX Cluster, Linux Cluster, and so forth) so that the cluster nodes are split between the sites. By having the automated failover/failback capabilities in the SAN, many application cluster technologies can leverage this capability to do "stretch" application clustering. Consult the application server specific information for information about setting up clustering to support this type of failover. In general, you must perform three tasks:

1. Set up an application server node in each site.
2. Configure the application server to connect to the volumes on the SAN.
3. If using multiple subnets in the storage cluster, log on to both VIPs in the cluster.

If you are using the HP P4000 DSM for MPIO, see "Using the HP P4000 DSM for MPIO," in the *HP StorageWorks P4000 Windows Solution Pack User Manual*.

Using the Primary Site designation in two-site configurations

The Multi-Site SAN software supports designating a particular site as Primary so that it does not go offline when the secondary site goes offline, or when the network link between sites goes offline. This is specific to the two-site configuration where, without this designation, there would be no way for the system to determine which site should remain online in the event of a network separation (classically known in the clustering/distributed systems world as a "split-brain" system).

To set up this type of configuration:

1. Designate one of the sites as primary.

2. Start a majority of managers in the primary site.

Failure modes and recovery steps

In general, there are three failure modes that a configuration like this can encounter that you should be aware of. Table 5 on page 19 summarizes the failure conditions and what you should do to rectify the issue.

Table 5 Failure conditions and resolutions

Failure condition	Resolution steps
Secondary site offline	None. The primary site will continue to operate.
Network link between sites fails	None. The primary site will continue to operate.
Primary site offline	1. If any of the storage nodes running managers in the primary site can be brought back online, you should do this first. 2. If all storage nodes in the primary site are lost forever, you can use the Recover Quorum command-line function to recover quorum in the secondary site manually. For more information, see "Recover quorum operation," next in this section

Recover quorum operation

The Recover Quorum function is introduced with the Multi-Site SAN as a way to bring volumes back online in the event that quorum in the management group is lost.

△ CAUTION:

The Recover Quorum function should only be used as a last resort. The failed site must be completely failed to use Recover Quorum safely on the remaining site.

The Recover Quorum command-line function performs the following actions:

1. Removes all offline managers from the management group.
2. Ejects all offline storage nodes from the management group.
3. The SAN/iQ software is modified to prevent any data integrity issues in the event any of the offline storage nodes are brought back online.

Best practices for configuring quorum and managers in a Multi-Site SAN environment

Configuring the appropriate number of managers in a Multi-Site SAN on the appropriate storage nodes is an important part of setting up the Multi-Site SAN for high availability and fault tolerance.

Table 6 on page 20 identifies common Multi-Site SAN configurations and the recommended manager configurations.

Table 6 Common Multi-Site SAN configurations with recommended managers

Multi-Site SAN configuration	Manager configuration	Number of managers and number Needed to maintain quorum
Two physical sites, one site is primary	For four storage nodes, run: • two managers in the primary site • one manager in the secondary site For six storage nodes or more, run: • three managers in the primary site • two managers in the secondary site	• Manager count = three Number for quorum = two • Manager count = five Number for quorum = three
Two physical sites, logical third site	For two storage nodes, run: • one manager in each site • one failover manager in the logical third site • two managers in each site, one failover manager in the logical third site For four storage nodes or more, run: • two managers in each site • one failover manager in the logical third site	• Manager count = three Number for quorum = two • Manager count = five Number for quorum = three
Three physical sites, cluster spans two sites	For two storage nodes, run: • one manager in each site • one failover manager in the third site For four storage nodes, run: • two managers in each site • one failover manager in the third site	• Manager count = three Number for quorum = two • Manager count = five Number for quorum = three
Three physical sites	For three storage nodes, run: • one manager in each site For six or more storage nodes, run: • two managers in each site	• Manager count = three Number for quorum = two • Manager count = six Number for quorum = four

2 Implementing a Multi-Site SAN

After you have designed the Multi-Site SAN configuration, you are ready to implement it.

Before you begin

Install Multi-Site SAN feature keys on the storage nodes you are going to use.

Register your storage nodes. For more information, see "Registering Advanced Features" in the *HP StorageWorks P4000 SAN Solution User Manual* for information about registering your storage nodes.

Have the following information ready:

- The layout for storage nodes and sites
- IP addresses, subnets, and VIPs to use for the Multi-Site cluster. You can use one VIP per subnet.
- (Optional) IP address and subnet for the Failover Manager

Preparing to install a new Multi-Site SAN

After you have finished planning the Multi-Site SAN and installed feature keys on the storage nodes for the SAN, you are ready to begin installation.

Preparing to convert an existing SAN to a Multi-Site SAN

Converting an existing SAN to a Multi-Site SAN includes planning the site layout and network information along with the following considerations:

- Do you need to move/remove data from the existing SAN before reconfiguring it?
- Do you need to purchase additional storage nodes to meet the requirement that each site have the same number of storage nodes?
- Do you need to reconfigure your network to take advantage of the Multi-Site SAN capabilities?

  See "Designing the network for the Multi-Site SAN" on page 13 for detailed information about setting up your network appropriately.

Installing the Multi-Site SAN

To install a Multi-Site SAN, use the Management Groups, Clusters, and Volumes Wizard.

Figure 12 Using the management group wizard to create a Multi-Site SAN

NOTE:

You must have at least one storage node for each site in the Available Nodes pool before starting.

Creating the management group

The first part of the wizard creates the management group, which includes configuring the optimal number of managers for the initial setup. After you have the Multi-Site clusters configured, you need to review the manager configuration and adjust it according to the Multi-Site configuration you have planned. See "Configuring managers" on page 25.

Creating Multi-Site clusters and volumes

The path through the Multi-Site cluster portion of the wizard follows the steps listed below. The steps are described in greater detail in"Using the Multi-Site cluster wizard" on page 22 .

1. Create a cluster.
2. Create a site and then assign storage nodes to the site.
3. Create the additional sites with storage nodes, up to a total of three sites.
4. Add one or more VIPs, depending on the number of subnets you have configured.
5. Create volumes using the data protection levels dictated by your Multi-Site cluster design.

Using the Multi-Site cluster wizard

After creating the management group, the Create a Cluster window of the wizard is displayed.

1. Select **Multi-Site Cluster** and then click **Next**.

   The Create Cluster window opens, displaying all of the storage nodes you designated for the management group.
2. Select all the storage nodes targeted for all of the sites you plan to create, type a name for the cluster, and then click **Next**.

Creating sites

When creating a new site, you first create the site and then assign storage nodes to that site:

1. Click **New** to open the New Site window.
2. Enter a name for the site and an optional description.
3. [Optional] Make the site primary, if appropriate.
4. Click **Add** in the Site Nodes section and then select the storage nodes for this site.
5. Click **OK** when you have finished creating the site.

   The Set up Sites window opens, displaying the first site created and the remaining unassigned storage nodes in the list.
6. Repeat these steps to create additional sites.
7. When you have finished creating sites, click **Next** to continue with the wizard.
8. Click **Add** to assign each VIP and Subnet Mask as planned for your site layout.

NOTE:

You can have one VIP per subnet.

9. Click **Next** to continue.
10. Click **Finish** when you have finished defining the volume.

    The management group, sites, Multi-Site cluster, and the volume are all created. This process takes some minutes.

Viewing the Multi-Site cluster information

You can view the completed sites and related information in two places. The Cluster tab view shows the sites in relation to the cluster.

- The Cluster tab view (Figure 13) which shows the sites.
- The Sites view (Figure 14) which shows detailed information about the storage nodes in the cluster.

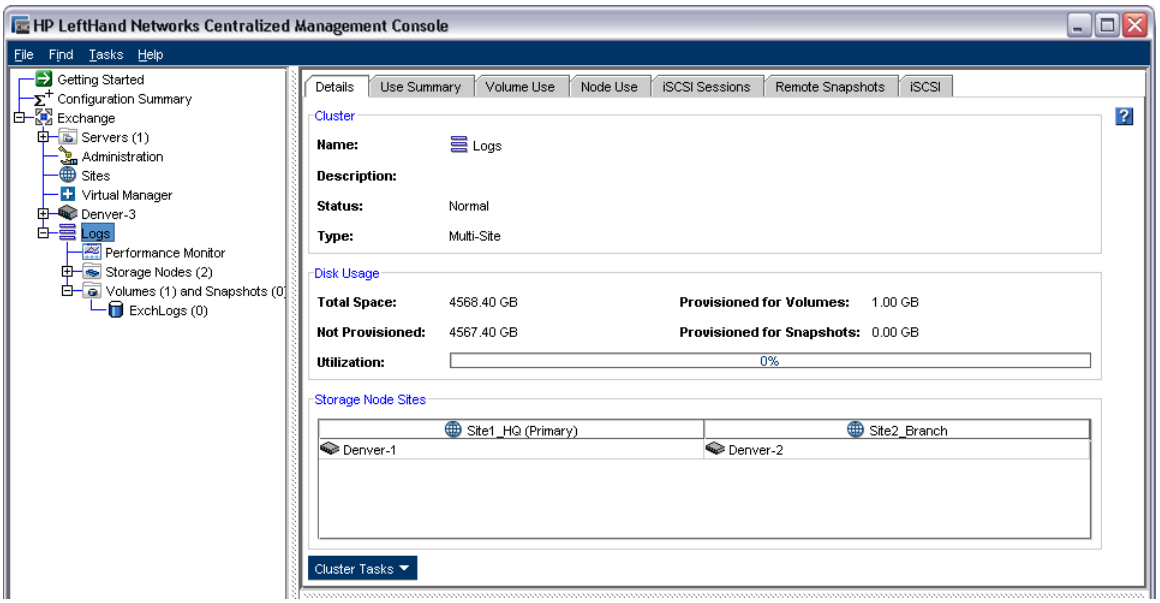

Figure 13 Viewing the sites in the cluster Details tab

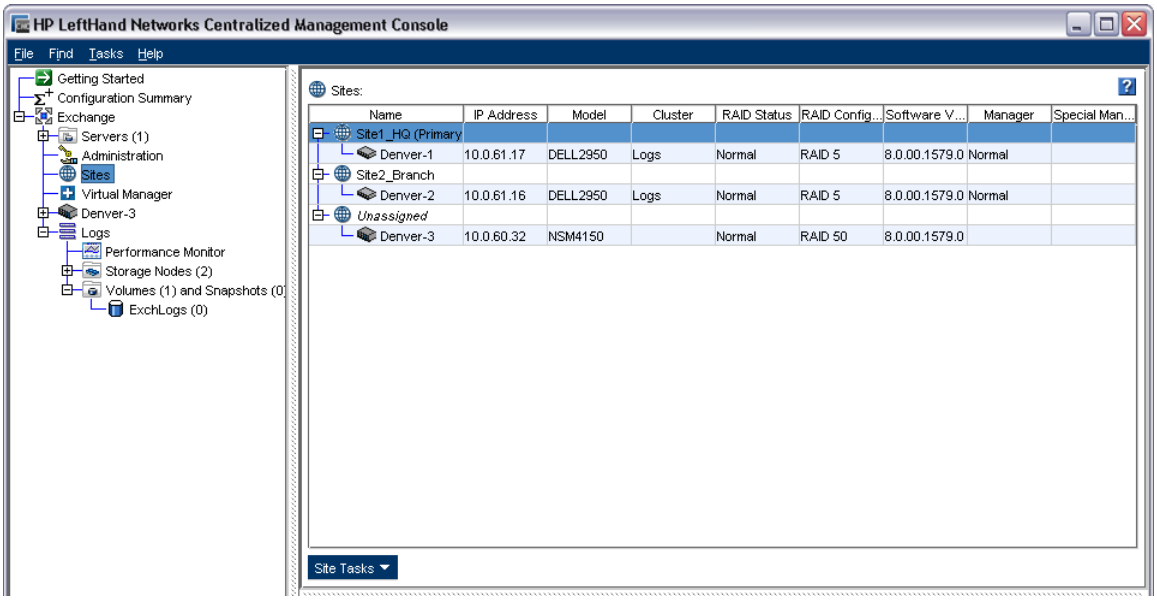

Figure 14 Viewing the Sites Details tab window

Converting an existing SAN to a Multi-Site SAN

Plan the conversion according to the criteria described in "Preparing to convert an existing SAN to a Multi-Site SAN" on page 21.

Prerequisites

- Upgrade all storage nodes to SAN/iQ software version 7.0 or later.
- Obtain the feature keys for storage nodes to be used in the Multi-Site cluster.
- Apply the feature keys to the storage nodes.

- Add or remove storage nodes to the cluster if necessary to ensure you can have equal numbers of storage nodes in each site. The wizard will not let you create sites with different numbers of storage nodes; nor will it let you leave unassigned any storage nodes currently in the cluster.

Using the Multi-Site cluster wizard

Use the Management Groups, Clusters, and Volumes wizard to convert an existing standard cluster to a Multi-Site cluster.

1. Start the wizard from the Getting Started Launch Pad.
2. In the wizard, select **Existing Management Group** and then click **Next**.
3. Select the management group that contains the cluster you want to convert.
4. Click **Next** to open the Create a Cluster window.
5. Select **Existing Cluster** and **Convert a Standard Cluster to a Multi-Site Cluster** and then click **Next**.
6. Click **Next** to select a specific cluster.
7. Select the cluster you want to convert and then click **Next**.
8. Follow the steps in "Creating sites" on page 23 to create the sites using the listed storage nodes.
9. After creating the sites, click **Next** The wizard opens the Assign Virtual IPs window.

   If the current cluster has a VIP assigned, it shows in this window. You can either use this VIP or assign new ones, according to the network plan for the Multi-Site cluster.
10. Assign one or more VIPs as planned for your site layout, and then click **Next**.

    The Create Volume window opens with the Skip Volume Creation box checked. If you want to create a new volume for your Multi-Site cluster, clear the box and fill in the volume information.
11. Click **Next** to complete the wizard.
12. Verify the settings for the sites you have just configured.
13. Click **Close** when you finish reviewing the Summary information.
14. Select the **Sites** node in the navigation window to view the new site configuration.

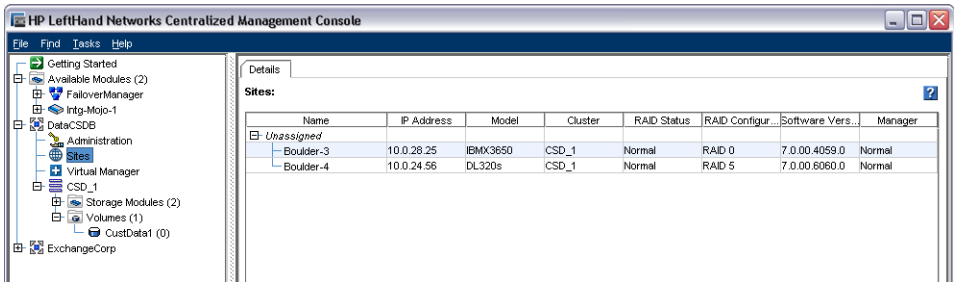

Figure 15 Sites Details table

Configuring managers

After you have set up your management group, cluster, and sites, you may need to change the manager configuration of your Multi-Site SAN. You also need to add the Failover Manager in all but one type of configuration.

Overview

The SAN/iQ software automatically starts the regular manager on each storage node when creating a management group. The number of managers started depends on the number of storage nodes in the group and whether you create a Multi-Site cluster. In some configurations, a Virtual Manager is also added to the management group. If, however, you plan to use a Failover Manager instead, you must manually install the Failover Manager on a separate server on the network and then add it to the management group as required by the design of your Multi-Site SAN. We recommend using the Failover Manager instead of the Virtual Manager because of its automated failover/failback capabilities.

Configuring regular managers

After you finish creating the management group, cluster, and sites for your Multi-Site SAN, review the manager configuration and change it as necessary. See Table 2 on page 10.

For information about management groups, managers, and starting and stopping managers, see the chapter "Working with Management Groups" in the *HP StorageWorks P4000 SAN Solution User Manual*.

Fault tolerant manager configurations

The SAN/iQ software tracks the configuration of managers to ensure that you have a fault-tolerant configuration. You may encounter messages with warnings or prerequisites you must meet before you can change the managers.

For example, if you create a management group and a Multi-Site cluster using two or four storage nodes, the management group has regular managers running on each storage node plus a Virtual Manager added to the management group. If you want to add a Failover Manager or stop a manager on one of the storage nodes, you must first delete the Virtual Manager from the management group. Then you can make the other changes.

Installing and configuring a Failover Manager

If you are using a Multi-Site SAN configuration that calls for a Failover Manager, you must install the Failover Manager from the SAN/iQ Management Software DVD onto a separate server on the network. The Failover Manager can run on VMware ESX Server, VMware Server, or VMware Player. For detailed instructions for installing and configuring the Failover Manager see "Using Specialized Managers" in the *HP StorageWorks P4000 SAN Solution User Manual*.

Best practice

Install the Failover Manager on a server in a separate physical location from the SAN sites, and then add it as a separate site in the management group. Doing so will prevent loss of quorum in a site failure, regardless of which single site fails.

Adding the Failover Manager to the management group and site

Adding the Failover Manager to the management group will likely cause the manager configuration to become less than optimal. Be sure to change the manager configuration to best reflect the new configuration of the management group. See the *HP StorageWorks P4000 SAN Solution User Manual*, Chapter 9, "Working with Management Groups," the section "Optimum Number of Managers."

1. Use **Find** to search for the Failover Manager, which appears in the Available Nodes pool in the CMC.

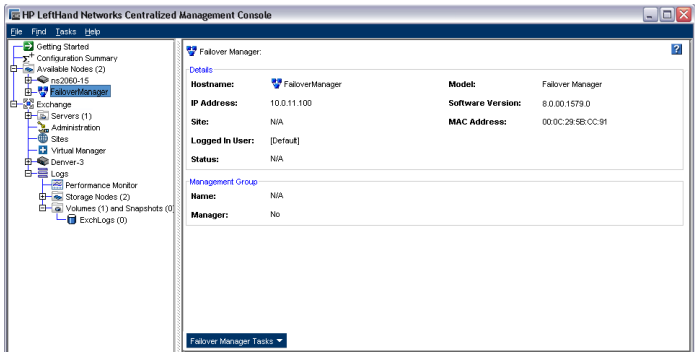

Figure 16 Finding the Failover Manager in the Available Nodes pool

2. Right-click on the **Failover Manager** and then select **Add to Existing Management Group**.
   - If there is more than one management group, select the correct one from the list that opens.
   - If there is only one management group, the Failover Manager is added to the group.
3. Select the **Sites** node in the navigation window.
4. Right-click and then select **New Site**.
5. In the New Site window, type the name and optional description for the Failover Manager site.
6. Click **Add** in the Site Nodes section.
7. In the Add Nodes to Sites window, select the **Failover Manager** from the list and then click **OK**.
8. Click **OK** on the New Site window.

   The new Failover Manager site appears in the Sites Details tab.

NOTE:

You can also add the Failover Manager to an existing site.

Configuring sites

A site is a user-designated location in which storage nodes are installed. Multi-Site SAN configurations have multiple sites with storage nodes in each site, and each site has its own subnet. A site can be a logical configuration, such as a subnet within the same data center, department, or application. Volumes are replicated across sites, which ensures high availability and redundancy.

Sites have the following characteristics:

- Have equal numbers of storage nodes, except for the Failover Manager, which can reside in a site by itself.
- Correlate to a geographical location and/or a configuration entity such as a subnet or power grid
- May be marked as primary
- Have a planned number of managers
- Have a planned data protection level

Viewing sites in the CMC

To view sites in the CMC, select the **Sites** node in the navigation window. The Sites node is displayed in every management group. Typically, however, the Sites Details tab shows the storage nodes as unassigned. The site designation has no effect outside of the Multi-Site SAN configuration.

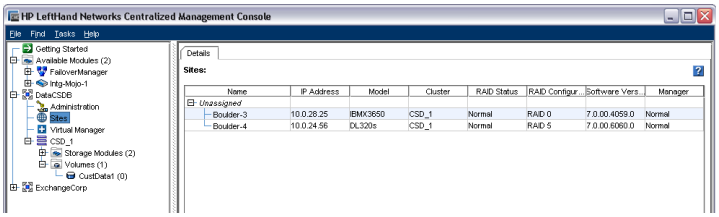

Figure 17 No sites in a standard cluster

When there is a Multi-Site cluster in the management group, the Sites Details tab shows the site assignments of the storage nodes.

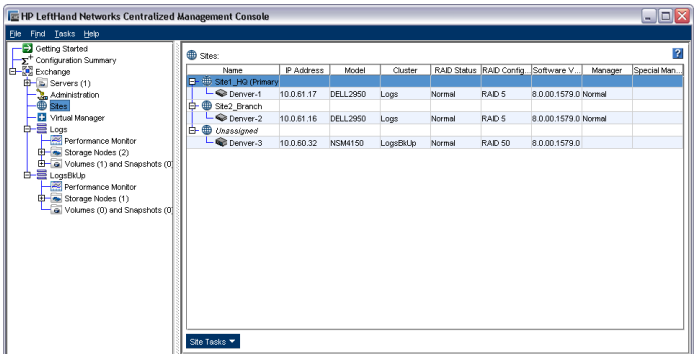

Figure 18 Site assignments in a Multi-Site cluster

A management group can contain multiple clusters, some of which are Multi-Site and some of which are not. The Sites tab shows which clusters have sites and which do not.

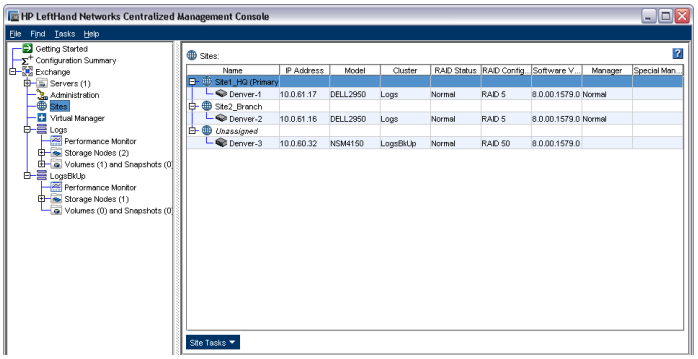

Figure 19 Management groups can have both standard and Multi-Site clusters

Designating a primary site

Use the Primary designation when you are using the Multi-Site cluster to maintain a primary site and a second site as a backup site. Designate the site as Primary when you create the site, or you can edit an existing site to make it primary. Only one site can be designated as Primary in a Multi-Site cluster. Make sure that the majority of managers are in the Primary site.

Adding a site

Create new sites directly in the CMC if you want to add a site to an existing cluster. You must have equal numbers of storage nodes in each site. The Failover Manager can occupy a site by itself.

1. Select the **Sites** node in the navigation window.
2. Right-click and then select **New Site**.
3. Type in a site name and optional description.
4. (Optional) If you are setting up a primary configuration, select the Primary box.
5. Click **Add** to add storage nodes to the site.
6. Select the appropriate nodes and then click **OK** when you are finished.
7. Click **OK** again when you have finished creating the new site.

   The Sites Details tab displays the new site with its storage node(s).

Editing a site

You can edit any aspect of the site, including:

- Changing the name and description of the site
- Changing the primary designation

NOTE:

Only one site can be primary.

- Adding or removing nodes.

IMPORTANT:

All sites must have the same number of storage nodes, except for the Failover Manager, which can reside in a site by itself.

To edit a site

1. Right-click the **Sites** node in the navigation window and then select **Edit Sites**.
2. Select the site you want to edit and then click **Continue**.
3. Make the desired changes.
4. Click **OK** when you are finished.

Deleting a site

You would delete a site from a Multi-Site cluster when you are either reconfiguring your SAN or recovering from a disaster or other failure.

Deleting a site changes the cluster from a Multi-Site cluster to a standard cluster. If you then remove the storage node(s) assigned to that site, data that resides on the cluster will restripe.

1. Select the **Sites** node in the navigation window.
2. Select the site you want to delete in the Details tab, right-click, and then select **Delete Site**.

NOTE:

If you delete a site that contains a storage node, the warning shown in Figure 20 on page 30 is displayed.

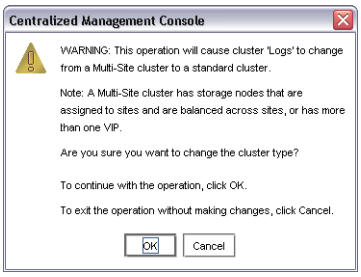

Figure 20 Deleting a site that contains a storage node causes this warning

3. Click **OK** to delete the site.

Adding storage nodes to a Multi-Site cluster

After your Multi-Site cluster is configured and operating, you may want to add storage nodes to the cluster for additional capacity or to create a new site. While you can add storage nodes to the Multi-Site cluster at any time, you should add them in a way that does not violate the requirements of the Multi-Site SAN.

See "Adding a site" on page 29 if you plan to add the storage nodes to a new site in the cluster.

Adding additional capacity to existing sites

Plan the following items according to your overall Multi-Site SAN configuration:

- Plan to add the same number of storage nodes to each site.
- Obtain feature keys for the new storage nodes.
- Plan the IP addresses and locations for the new storage nodes to fit into the overall design for your Multi-Site SAN.

To add capacity

1. Add the license keys to the new storage nodes.

   See "Registering Advanced Features" in the *HP StorageWorks P4000 SAN Solution User Manual* for more information about license keys.
2. Add the new storage nodes to the management group.
3. Select the management group **Sites** node in the navigation window.
4. Select the site to which you want to add a storage node.
5. Right-click and then select **Edit Site**.

6. Click **Add** to add the storage node.

   The Add Nodes to Site window opens.
7. Select the storage node from the list and then click **OK**.
8. Click **OK** to finish adding the storage node to the site.
9. Repeat these steps for all the storage nodes you are adding.
10. Add the storage nodes to the Multi-Site cluster, using the Edit Cluster window.

Removing storage nodes from a Multi-Site cluster

Removing storage nodes from a Multi-Site cluster affects the capacity of the cluster and the Multi-Site configuration.

- Ensure that the capacity remaining after you remove the storage node(s) is sufficient for the data on the volumes.
- To maintain the Multi-Site configuration, plan to remove equal numbers of storage nodes from each site.

To remove the storage nodes from the site

1. Select the **Sites** node in the navigation window.
2. Select the site from which you want to remove the storage node.
3. Right-click and then select **Edit Site**.
4. Select the storage node from the list and then click **Remove**.

   The storage node is removed from the list.
5. Click **OK** on the Edit Site window.

   A message opens, warning that the Multi-Site cluster will change to a standard cluster if you remove the storage node.
6. Click **OK** to continue.
7. Repeat these steps to remove storage nodes from additional sites.

NOTE:

Adding the Failover Manager will likely cause the manager configuration to become less than optimal. Be sure to change the manager configuration to best reflect the new configuration of the management group. See the section "Optimum Number of Managers" in the chapter, "Working with Management Groups" in the *HP StorageWorks P4000 SAN Solution User Manual*.

3 Support and other resources

Contacting HP

For worldwide technical support information, see the HP support website:

http://www.hp.com/support

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

http://www.hp.com/go/e-updates

After registering, you will receive email notification of product enhancements, new driver versions, firmware updates, and other product resources.

New and changed information in this edition

The following additions and changes have been made for this edition:

- The following information has been updated:
  - New data protection levels, called Network RAID
- A new Support and Other Resources chapter has been added.
- The Preface was removed

Related information

The following documents [and websites] provide related information:

- *HP StorageWorks P4000 SAN Solution User Manual*

You can find these documents from the Manuals page of the HP Business Support Center website:

http://www.hp.com/support/manuals

In the Storage section, click **Disk Storage Systems** and then select P4000 SAN Solutions.

HP websites

For additional information, see the following HP websites:

- http://www.hp.com
- http://www.hp.com/go/storage
- http://www.hp.com/service_locator
- http://www.hp.com/support/manuals
- http://www.hp.com/support/downloads

Typographic conventions

Table 7 Document conventions

Convention	Element
Blue text: Table 7	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	Website addresses
Bold text	• Keys that are pressed • Text typed into a GUI element, such as a box • GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes
Italic text	Text emphasis
`Monospace` text	• File and directory names • System output • Code • Commands, their arguments, and argument values
`Monospace, italic` text	• Code variables • Command variables
`Monospace, bold` text	Emphasized monospace text

IMPORTANT:

Provides clarifying information or specific instructions.

NOTE:

Provides additional information.

TIP:

Provides helpful hints and shortcuts.

HP product documentation survey

Are you the person who installs, maintains, or uses this HP storage product? If so, we would like to know more about your experience using the product documentation. If not, please pass this notice to the person who is responsible for these activities.

Our goal is to provide you with documentation that makes our storage hardware and software products easy to install, operate, and maintain. Your feedback is invaluable in letting us know how we can improve your experience with HP documentation.

Please take 10 minutes to visit the following web site and complete our online survey. This will provide us with valuable information that we will use to improve your experience in the future.

http://www.hp.com/support/storagedocsurvey

Thank you for your time and your investment in HP storage products.

Glossary

The following glossary provides definitions of terms used in the SAN/iQ software and the HP LeftHand SAN Solution.

active monitoring Active monitoring tracks the health of the storage node using notifications such as emails, alerts in the CMC, and SNMP traps.

acting primary volume The remote volume, when it assumes the role of the primary volume in a failover scenario.

Active-Passive A type of network bonding which, in the event of a NIC failure, causes the logical interface to use another NIC in the bond until the preferred NIC resumes operation. At that point, data transfer resumes on the preferred NIC.

Adaptive Load Balancing A type of network bonding in which the logical interface performs load balancing of data transmission.

add-on application An additional feature purchased separately from the SAN/iQ software.

application-managed snapshot Snapshot of a volume that is taken while the application that is serving that volume is quiesced. Because the application is quiesced, the data in the snapshot is consistent with the application's view of the data. That is, no data was in flight or cached waiting to be written.

authentication group For release 7.0 and earlier, identifies the client or entity accessing the volume. Not used in release 8.0 and later.

Auto Discover A feature in the CMC that automatically searches for storage nodes on the subnet the CMC is connected to. Any storage nodes it discovers appear in the navigation window on the left side of the CMC.

Bond0 Interface created for network interface failover and only appears after configuring for failover.

bonding Combining physical network interfaces into a single logical interface.

boot device Compact flash cards from which the storage node boots up. Also known as disk-on-modules or DOMs.

CHAP Challenge-Handshake Authentication Protocol (CHAP) is a standard authentication protocol.

clone point The snapshot that has two or more volumes associated with it. A clone point is created when a SmartClone volume is created from a snapshot or from snapshot temporary space.

CLI Command-line interface for the SAN/iQ software.

cluster A cluster is a grouping of storage nodes that create the storage pool from which you create volumes.

CMC	Centralized Management Console. See HP LeftHand Centralized Management Console.
communication mode	The unicast communication among storage nodes and application servers.
community string	The community string acts as an authentication password. It identifies hosts that are allowed read-only access to the SNMP data.
Configuration Summary	The Configuration Summary displays an overview of the volumes, snapshots, storage nodes, and iSCSI sessions in the HP LeftHand Storage Solution. It provides an overview of the storage network broken out by management groups.
data center	Also known as a "Site." A data center is a physical location in your environment where application servers, SAN storage and network equipment reside. In the SAN/iQ Multi-Site software, a data center is typically referred to as a site.
disaster recovery site	Similar to a secondary site, the disaster recovery site is used to operate the SAN in the event of a disaster.
disk status	Whether the disk is • Active - on and participating in RAID • Uninitialized or Inactive - On but not participating in RAID • Off or Missing - Not on • DMA Off - disk unavailable due to faulty hardware or improperly seated in the chassis
DSM	Device Specific Module
DSM for MPIO	The HP LeftHand Networks vendor-specific DSM that interfaces with the Microsoft MPIO framework.
failback	After failover, the process by which you restore the primary volume and turn the acting primary back into a remote volume.
failover	The process by which the user transfers operation of the application server over to the remote volume. This can be a manual operation, a scripted operation, or VMware enabled.
Failover Manager	A specialized manager running as a VMware appliance that allows you to place a quorum tie-breaker node into a 3rd location in the network to provide for automated failover/failback of the Multi-Site SAN clusters. The Failover Manager is designed to run on VMware ESX Server, VMware Server and VMware Player. It is installed on hardware separate from the SAN hardware.
failover recovery	After failover, the process by which the user chooses to fail back to the primary volume or to make the acting primary into a permanent primary volume.
frame size	The frame size specifies the size of data packets that are transferred over the network.
full provisioning	Full provisioning reserves the same amount of space on the SAN as is presented to application servers.

ghost storage node	When using Repair Storage Node, a “ghost” storage node acts as a placeholder in the cluster, keeping the cluster intact, while you repair or replace the storage node.
Graphical Legend	Describes all the icons used in the CMC • Items tab - displays the icons used to represent virtual items displayed in the CMC • Hardware tab - displays the icons that represent the physical storage units.
hardware reports	Hardware reports display point-in-time statistics about the performance and health of the storage node, its drives, and configuration.
hostname	The hostname on a storage node is the user-definable name that displays below the storage node icon in the network window. It is also visible when the users browse the network.
HP LeftHand Centralized Management Console	Management interface for the SAN/iQ software.
ID LED	LED lights on the physical storage node so that you can find that node in a rack [NSM 260 only].
iSCSI	Internet SCSI. The iSCSI protocol defines the rules and processes for transporting SCSI (block-level) data over a TCP/IP network.
iSCSI load balancing	Improves iSCSI performance and scalability by distributing iSCSI sessions for different volumes evenly across storage nodes in a cluster.
license keys	A license key registers a storage node for add-on applications. Each storage node requires its own license key.
Link Aggregation Dynamic Mode	A type of network bonding in which the logical interface uses both NICs simultaneously for data transfer.
log files	Log files for the storage node are stored both locally on the storage node and are also written to a remote log server.
logical site	This site is on an isolated network and power connection than the other sites. However, it can be in the same physical location as one of the real sites. Also, a site for a Failover Manager.
management group	A collection of one or more storage nodes which serves as the container within which you cluster storage nodes and create volumes for storage.
managers	Manager software runs on storage nodes within a management group. You start managers on designated storage nodes to govern the activity of all of the storage nodes in the group.
MIB	The Management Information Base provides SNMP read-only access to the storage node.
monitored variables	Variables that report health status of the storage node. These variables can be monitored using alerts, emails, and SNMP traps.

Multi-Site cluster A cluster of storage that spans multiple sites (up to 3). A Multi-Site cluster must meet at least one of the following conditions:

- Contain storage nodes that reside in two or more sites
- Contain storage nodes that span subnets
- Contain multiple VIPs. The cluster can have a single site, and the multiple VIPs make it a multi-site cluster.

network RAID Synchronous replication, mirroring or parity protection on a volume-by-volume basis. Protecting data for a volume across all storage nodes in the cluster. Network RAID-10, 10+1 or 10+2 is required to protect data in an HP LeftHand SAN solution.

network window Graphically depicts the status of each storage node. Storage nodes on the network are either available or part of a management group.

NTP Network Time Protocol

parity In RAID5, redundant information is stored as parity distributed across the disks. Parity allows the storage node to use more disk capacity for data storage.

peer site Absence of a primary site designation makes all the sites peer sites.

primary site A site designation assigned by the administrator in the HP LeftHand Centralized Management Console. A primary site is more important than a secondary site. In this setup, you would run a majority of managers in the primary site. In a two-site setup, this allows the primary site to stay online even if the network link between the primary and secondary sites fails. Typically, the primary site has majority/all of the application servers. In configurations that do not designate a primary site, the sites are referred to as “peer” sites.

original primary volume The primary volume that fails and then is returned to service.

overprovisioned cluster An overprovisioned cluster occurs when the total provisioned space of all volumes and snapshots is greater than the physical space available on the cluster. This can occur when there are snapshot schedules and/or thinly provisioned volumes associated with the cluster.

point-in-time snapshot Snapshots that are taken at a specific point in time, but an application writing to that volume may not be quiesced. Thus, data may be in flight or cached and the actual data on the volume may not be consistent with the application's view of the data.

preferred interface A preferred interface is the interface within an active backup bond that is used for data transfer during normal operation.

primary snapshot A snapshot of the primary volume which is created in the process of creating a remote snapshot. The primary snapshot is located on the same cluster as the primary volume.

primary volume The volume which is being accessed (read/write) by the application server. The primary volume is the volume that is backed up with Remote Copy.

quorum A majority of managers required to be running and communicating with each other in order for the SAN/iQ software to function.

RAID device RAID (originally redundant array of inexpensive disks, now redundant array of independent disks) refers to a data storage scheme using multiple hard drives to share or replicate data among the drives.

RAID levels Type of RAID configuration.

- RAID0 - data striped across disk set
- RAID1 - data mirrored from one disk onto a second disk
- RAID10 - mirrored sets of RAID1 disks
- RAID5 - data blocks are distributed across all disks in a RAID set. Redundant information is stored as parity distributed across the disks.
- RAID50 - mirrored sets of RAID5 disks.

RAID quorum Number of intact disks required to maintain data integrity in a RAID set.

RAID rebuild rate The rate at which the RAID configuration rebuilds if a disk is replaced.

RAID status Condition of RAID on the storage node

- Normal - RAID is synchronized and running. No action is required.
- Rebuild - A new disk has been inserted in a drive bay and RAID is currently rebuilding. No action is required.
- Degraded - RAID is not functioning properly. Either a disk needs to be replaced or a replacement disk has been inserted in a drive.
- Off - Data cannot be stored on the storage node. The storage node is offline and flashes red in the network window.

register Register individual storage nodes to use add-on applications. Registration requires sending in the storage node serial numbers to purchase the license keys, which are then applied to the storage node.

remote copy pair The primary volume and its associated remote volume.

remote snapshot An identical copy of a primary snapshot. The remote snapshot is located on the same cluster as the remote volume.

remote volume The volume that resides in the Remote Copy location where the remote snapshots are created. The remote volume contains no data. It acts as a pointer to tell the system where to make the copy of the primary snapshot. The remote volume can be stored in these ways:

- In the same cluster in the same management group
- In a different cluster in a different management group
- In a different cluster in the same management group

Repair storage node Creates a placeholder in the cluster, in the form of a “ghost” storage node, that keeps the cluster intact while you remove the storage node to replace a disk or replace the storage node itself, and return it to the cluster.

replication level In Release 8.5 this changes to data protection level. Prior to release 8.5, replication level is the term that designated how many copies of data to keep in the cluster.

replication priority Removed in Release 8.5. Prior to Release 8.5, replication priority allowed you to designate whether data availability or redundancy is more important in your

	configuration. Release 8.5 forward defaults to availability. This default can be changed using the Cliq Command Line Interface.
restripe	Striped data is stored across all disks in the cluster. You might change the configuration of a volume, for example, change data protection level, add a storage node, or remove a storage node. Because of your change, the pages in the volume must be reorganized across the new configuration. The system can keep track of several configuration changes at once. This means you can change configurations, even while a volume is in the midst of a different reconfiguration. In particular, if a reconfiguration was done by accident, you don't have to wait until it finishes to change back to the original configuration. See “Stripe”.
resync	When a storage node goes down, and writes continue to a second storage node, and the original store comes back up, the original storage node needs to recoup the exact data captured by the second storage node.
rolling back	Replaces the original volume with a read/write copy of a selected snapshot. New for release 8.0: The new volume retains the same name.
SAN/iQ interface	When you initially set up a storage node using the Configuration Interface, the first interface that you configure becomes the interface used for the SAN/iQ software communication.
secondary site	A site that is less important than the primary site. In this setup a minority of managers runs in the secondary site. In a two-site setup, this allows the secondary site to go offline if the network link between the Primary and secondary sites fails. Typically, the secondary site has a minority, or none, of the application servers. If the primary site fails, customers can manually recover quorum in the secondary site.
server	An application server that you set up in a management group and then assign volumes to it to provide access to those volumes.
shared snapshot	Shared snapshots occur when a clone point is created from a newer snapshot that has older snapshots below it in the tree. All the volumes created from the clone point will display these older snapshots that they share, as well as the clone point.
site	A user-designated location in which storage nodes are installed. Multi-Site SAN configurations have multiple sites with storage nodes in each site, and each site has its own subnet. A site can be a logical configuration, such as a subnet within the same data center, department, or application.
SmartClone volume	SmartClone volumes are space-efficient copies of existing volumes or snapshots. They appear as multiple volumes that share a common snapshot, called a clone point. They share this snapshot data on the SAN.
snapshot	A fixed version of a volume for use with backup and other applications.
snapshot set	Application-managed snapshots created for a volume set.
SNMP traps	Use traps to have an SNMP tool send alerts when a monitoring threshold is reached.
solution pack	HP LeftHand P4000 Windows Solution Pack

split mirror	A split mirror is a remote snapshot whose relationship to the primary volume has been severed. Split mirrors are usually created for one-time use and then discarded.
standard cluster	Also known as a “cluster.” A standard cluster is one that does not use any of the Multi-Site features within the SAN/iQ software. Standard clusters • Cannot contain storage nodes that are designated to reside in a site. • Cannot contain storage nodes that span subnets • Can only have a single VIP.
storage server	Storage server software maintains the customer's data. It reads to and writes from disks in response to customer reads and writes of SANiQ volumes.
stripe	Striped data is stored across all disks in the array, which increases performance but does not provide fault tolerance.
synchronize	The process of copying the most recent snapshot from the primary volume to a new remote snapshot. On failback, synchronization is the process of copying the most recent remote snapshot back to the primary volume. The CMC displays the progress of this synchronization. Also, you can manually synchronize if necessary to include data that is on the remote volume but not the primary.
target secret	Target secret is used in both 1-way and 2-way CHAP when the target (volume) challenges the iSCSI initiator.
temporary space	Temporary space is created when a snapshot is mounted for use by applications and operating systems that need to write to the snapshot when they access it. Temporary space can be converted to a volume using the SmartClone process.
thin provisioning	Thin provisioning reserves less space on the SAN than is presented to application servers.
Trap Community String	The Trap Community String is used for client-side authentication when using SNMP.
unicast	Communication between a single sender and a single receiver over a network.
VIP	virtual IP address
virtual IP address	A highly available address that ensures that if a storage node in a cluster becomes unavailable, servers can still access the volume through the other storage nodes in the cluster.
virtual machine	A virtual storage appliance that provides one or more simultaneous storage environments in which SAN/iQ may execute as though they were running on the bare iron.
virtual manager	A manager that is added to a management group but is not started on a storage node until it is needed to regain quorum.
volume	A logical entity that is made up of storage on one or more storage nodes. It can be used as raw data storage or it can be formatted with a file system and used by a host or file server.
volume set	Two or more volumes used by an application. For example, you may set up Exchange to use two volumes to support a StorageGroup: one for mailbox data and one for logs. Those two volumes make a volume set.

volume lists For release 7.0 and earlier, provide the link between designated volumes and the authentication groups that can access those volumes. Not used in release 8.0 and later.

volume size The size of the virtual device communicated to the operating system and the applications.

VSS Volume Shadow Copy Service

VSS Provider HP LeftHand P4000 VSS Provider is the hardware provider that supports the Volume Shadow Copy Service on the HP LeftHand Storage Solution.

writable space *See* temporary space

Index