# HP StorageWorks

# P4000 Remote Copy user guide


## Abstract

This guide provides information about configuring and using asynchronous replication of storage volumes and snapshots across geographic distances.

For the latest version of this guide, see the HP website http://www.hp.com/support/manuals.



Part Number: AX696-96032
Fourth edition: November 2010


**Legal and notice information**

# Contents

# Figures

# Tables

# 1 Understanding and planning Remote Copy

Remote Copy provides a powerful and flexible method for reproducing data and keeping that replicated data available for disaster recovery, business continuance, backup and recovery, data migration, and data mining.

## How Remote Copy works

Remote Copy uses the existing volume and snapshot features with replication across geographic distances to create remote snapshots. The geographic distance can be local (in the same data center or on the same campus), metro (in the same city), or long distance (cross-country, global).

For example, the accounting department in the corporate headquarters in Chicago runs the corporate accounting application and stores the resulting data. The designated backup site is in Detroit. Nightly at 11:00 p.m., accounting updates are copied to the Detroit backup facility using Remote Copy. Figure 1 on page 9 shows the basic flow of Remote Copy.

Reproducing data using Remote Copy follows a three-step process:

1. At the production location, create a snapshot of the primary volume. This is called the primary snapshot.
2. Create a remote volume at the remote location, and then create a remote copy of the primary snapshot to the remote volume.

   The system copies data from the primary snapshot to the remote snapshot.

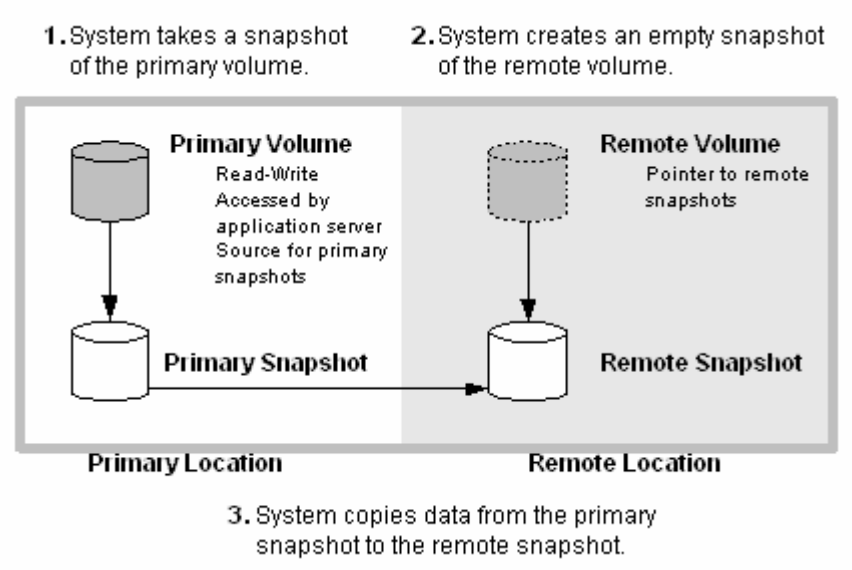


**Figure 1 Basic flow of Remote Copy**

NOTE:

Both primary and completed remote snapshots are the same as regular snapshots.

Remote Copy can be used on the same site, even in the same management group and cluster.

# Graphical representations of Remote Copy

The HP StorageWorks P4000 Centralized Management Console displays special graphical representations of Remote Copy.

## Copying the primary snapshot to the remote snapshot

When the primary snapshot is copying to the remote snapshot, the CMC depicts the process with a moving graphic of pages from the primary to the remote snapshot. The pages move in the direction of the data flow from primary to remote snapshot.

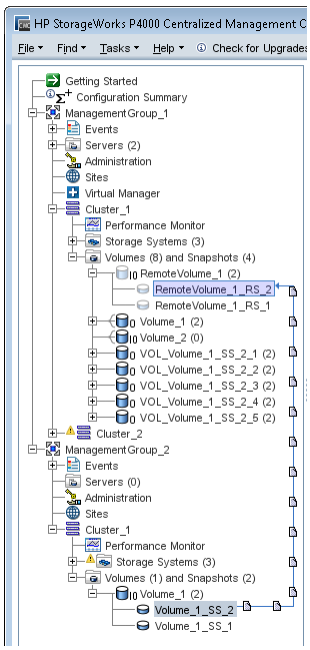


**Figure 2 Icons depicting the primary snapshot copying to the remote snapshot**

## Graphical legend for Remote Copy icons

The Graphical Legend window available from the Help menu depicts the icons associated with Remote Copy.

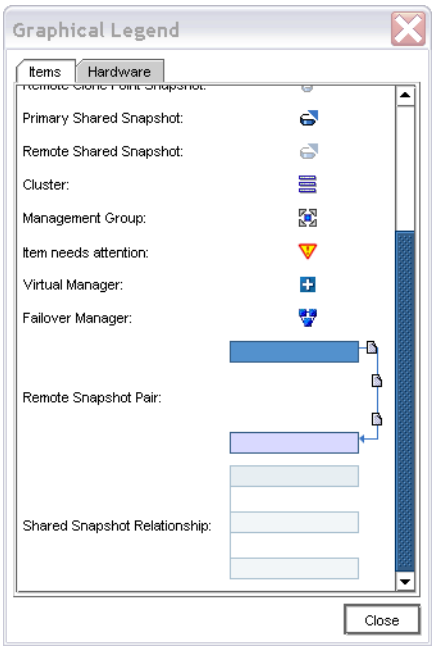


**Figure 3 Icons for Remote Copy in the Graphical Legend window**

## Remote Copy and volume replication

Remote Copy is an asynchronous replication of data. Network RAID is a synchronous replication of data. Using synchronous Network RAID on multiple storage systems within a cluster in combination with asynchronous Remote Copy on a different cluster of storage systems creates a robust, high-availability configuration.

## Uses for Remote Copy

Common uses for the Remote Copy application.

**Table 1 Uses for Remote Copy**

| Use Remote Copy for | How it works |
|---|---|
| Business continuance and disaster recovery | Remote Copy stores remote snapshots on a machine in a geographically separate location. The remote snapshots remain available in the event of a site or system failure at the primary site. |
| Off-site backup and recovery | Remote Copy eliminates the backup window on an application server by creating remote snapshots on a backup server, either local or remote, and back up from that server. |
| Split mirror, data migration, content distribution | Remote Copy makes a complete copy of one or more volumes without interrupting access to the original volumes. You can move the copy of the volume to the location where it is needed. |
| Volume clone | Remote Copy creates copies of the original volume for use by other application servers. |

## Benefits of Remote Copy

- Remote Copy maintains the primary volume's availability to application servers. Snapshots on the primary volume are taken instantaneously, and are then copied to remote snapshots in the off-site location.
- Remote Copy operates at the block level, moving large amounts of data much more quickly than file system copying.
- Snapshots are incremental, that is, snapshots save only those changes in the volume since the last snapshot was created. Therefore, failback may need to resynchronize only the latest changes rather than the entire volume.
- Remote Copy is robust. If the network link goes offline during the process, copying resumes where it left off when the link is restored.

# Planning for Remote Copy

Remote Copy works at the management group, cluster, volume, and snapshot levels. Review Table 2 on page 12 for common configurations at these levels.

**Table 2 Remote Copy, SAN/iQ, and storage systems**

| Storage system level | Remote Copy configuration |
|---|---|
| Management groups | • Create remote snapshots in the same management group or in a different management group than the primary volume.<br>• If using different management groups, the remote bandwidth setting of the management group containing the remote volume determines the maximum rate of data transfer to the remote snapshot.<br>• Event notification tells you when copies complete or fail. It also notifies you if a remote volume or snapshot is made primary or if the status of the connection between management groups containing primary and remote volumes changes. |
| Clusters | • Create remote snapshots in the same cluster or in a cluster different from the primary volume. |
| Volumes | • Primary volumes contain the data to be copied to the remote snapshot.<br>• Data is copied to the remote snapshot via the remote volume.<br>• The remote volume is a pointer to the remote snapshot. The remote volume has a size of 0 bytes. |
| Snapshots | • After data is copied from the primary snapshot to the remote snapshot, the remote snapshot behaves as a regular snapshot. |

## Planning the remote snapshot

### Prerequisites for creating a remote snapshot

- Log in to both the management group that contains the primary volume and the management group that contains the target cluster where the remote snapshot will be created.
- Designate or create a remote volume in that remote management group.
- Ensure there is enough space on the target cluster for the remote snapshot.

## Logging in to primary and remote management groups

Log in to both the primary and the remote management groups before you begin, or you must log in to the remote management group while creating a remote copy.

## Designating or creating the remote volume

Create a remote volume by using any of the following methods:

- Make an existing volume into a remote volume.
- Create a new remote volume while creating a remote snapshot.
- Create a new volume from the cluster Details tab window, and then select the Remote radio button on the Advanced tab of the New Volume window.

  From the menu bar, select **Tasks** > **Volume** > **New Volume**.

For more information about these methods of creating remote volumes, see "Creating a remote volume" on page 21.

# Using schedules for Remote Copy

Scheduled remote snapshots provide fault tolerance for business continuance and disaster recovery, and a consistent, predictable update of data for remote backup and recovery.

## Planning the Remote Copy schedule

Planning is critical. The following issues impact the amount of storage available in the system:

- Recurrence
- Capacity
- Retention

### Recurrence

How often do you want the snapshots created? The recurrence frequency must account for the amount of time it takes to complete a remote snapshot. For example, if your recurrence schedule is set for a new snapshot every four hours, you should ensure that the time to copy that snapshot to the remote location is less than four hours.

#### Testing the copy time

One way to check the time required to copy a snapshot is to run a test of the actual process. To test this process, take two remote snapshots of the primary volume. Because the first remote snapshot copies the entire volume, it takes longer to copy. The second remote snapshot copies only *changes* made to the volume since the first remote snapshot was taken. You create the second remote snapshot after the time interval you intend to schedule. Because of that, the copy time for the second remote snapshot is more representative of the actual time required for copying subsequent remote snapshots.

To test the copy time:

1. Create a remote snapshot of the primary volume.
2. Wait for the copy to finish.
3. Create another remote snapshot of the primary volume.

4. Track the time required to complete the second remote snapshot.

NOTE:

This is the minimum amount of time that you should allow between scheduled copies.

5. Check the remote bandwidth setting for the other management group by using the `Edit Management Group` command. This setting affects the time required to copy a remote snapshot.

## Capacity

Does the cluster that contains the remote volume have sufficient space to accommodate scheduled snapshots?

If the cluster does not have sufficient space available, the remote snapshot appears in the CMC and flashes red. On the Details tab of the remote snapshot, the status message displayed is `Read only, not enough space in cluster to start copy.`

## Retention policies

How long do you want to retain the primary snapshots? The remote snapshots? You can set different retention policies for the primary and remote snapshots. For example, you can choose to retain two primary snapshots and five remote snapshots. The number of snapshots retained refers to completed snapshots. Take the following characteristics of scheduled remote snapshots into account when planning retention policies.

- The SAN/iQ software never deletes the last fully synchronized remote snapshot.

  Under some circumstances, such as unpredictable network speeds or varying snapshot size, a scheduled remote snapshot may create primary snapshots so frequently that the remote copy process cannot keep up with them. The retention policies for scheduled remote snapshots ensure that such factors do not cause primary and remote snapshots to become unsynchronized. Regardless of the retention policy defined for scheduled remote snapshots, up to two additional snapshots may be retained by the system at any given time. These two additional snapshots include the snapshot that is in the process of being copied and the last fully synchronized snapshot. A fully synchronized snapshot is one that has completed copying so that the remote snapshot matches its corresponding primary snapshot.

- Up to two additional snapshots may be retained at any given time.

  Because the SAN/iQ software never deletes the last fully synchronized snapshot, a remote copy schedule may retain *n*+2 copies for a retention policy of *n* (the currently copying remote snapshot plus the last fully synchronized snapshot). Using the example above, if you have a retention policy for your remote copy schedule of two primary and five remote snapshots, the software may retain up to four primary and seven remote snapshots for a period of time. Table 3 on page 14 shows the maximum retained snapshots with respect to a specific retention policy.

**Table 3 Snapshot retention policy and maximum number of retained snapshots**

| Scheduled remote snapshot retention policy | Maximum number of snapshots retained |
|---|---|
| *n* of primary snapshots *x* of remote snapshots | *n* + 2 of primary snapshots *x* + 2 of remote snapshots |
| *n* of hours for primary snapshots *x* of hours for remote snapshots | *n* + 2 primary snapshots older than *nx* + 2 remote snapshots older than *x* |

| Scheduled remote snapshot retention policy | Maximum number of snapshots retained |
|---|---|
| *n* of days for primary snapshots *x* of days for remote snapshots | *n* + 2 primary snapshots older than *nx* + 2 remote snapshots older than *x* |
| *n* of weeks for primary snapshots *x* of weeks for remote snapshots | *n* + 2 primary snapshots older than *nx* + 2 remote snapshots older than *x* |

- A remote snapshot is deleted only after its corresponding primary snapshot is deleted.

  Additionally, a remote snapshot is deleted only after its counterpart primary snapshot. You cannot retain fewer remote snapshots than primary snapshots when setting your retention policies.

NOTE:

Over the course of time, through deletion of primary snapshots, if you accumulate more remote snapshots than primary snapshots, the remote snapshots become regular snapshots when their corresponding primary snapshots are deleted. You can identify them as remote snapshots by their names, since the naming convention is established as part of creating the remote snapshot schedule.

## Best practices

- Retain at least two primary snapshots to ensure that only incremental copying is required for remote snapshots.
- Review your remote copy schedule to ensure that the frequency of the remote copies correlates to the amount of time required to complete a copy.

Use the checklist in Table 4 on page 15 to help plan scheduled remote snapshots.

**Table 4 Scheduled Remote Copy planning checklist**

| Configuration category | Characteristic |
|---|---|
| **Scheduled snapshot** | |
| Start time | • Start date (mm/dd/yyyy) for the schedule to begin<br>• Start time (mm:hh:ss) for the schedule to begin |
| Recurrence | • Recurrence is a yes or no choice. Do you want to take a remote snapshot one time in the future and not have it recur, or do you want a remote snapshot to be taken on a regular schedule?<br>• Frequency (minutes, hours, days, or weeks) determines the interval between recurring, scheduled, remote snapshots. |
| **Primary setup** | |
| Retention | Select one of the following options:<br>• Maximum number of snapshots (#)<br>• Set period of time (minutes, hours, days, or weeks) |
| **Remote setup** | |
| Management group | The management group that contains the remote snapshot |
| Volume | The remote volume for the remote snapshots |

| Configuration category | Characteristic |
| --- | --- |
| Retention | Select one of the following options:<br>• Maximum number of snapshots. This number equals completed snapshots only. In-progress snapshots take additional space on the cluster while copying. Also, the system will not delete the last fully synchronized snapshot. For space calculations, figure *n*+2 where *n*=maximum number of snapshots.<br>• Set period of time (minutes, hours, days or weeks) |

# 2 Using Remote Copy

For information about how Remote Copy works and how to plan capacity for Remote Copy, see "Understanding and planning remote copy" on page 9.

## Working with remote snapshots

Remote snapshots are a core component of Remote Copy. Remote Copy uses the existing volume and snapshot capabilities to replicate or copy, the data across geographic distances.

### New for release 8.5 and greater

You can create application-managed snapshots that use VSS to quiesce the application before creating the snapshots. 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 when the application created the snapshot.

### Creating a remote snapshot

Creating a remote snapshot is the main task when working with Remote Copy. You can either create a one-time remote snapshot or set up a schedule for recurring remote snapshots. Many of the characteristics for either case are the same.

Create a remote snapshot by using the following steps:

- Log in to the primary management group.
- Log in to the remote management group.
- Create a primary snapshot of the primary volume manually. When setting up a schedule to create a remote snapshot of a volume, the software automatically creates a primary snapshot, which is then copied to the remote volume.
- Either create a remote volume on a remote management group, or select an existing remote volume.
- Create the remote snapshot.

#### Best practice

The best way to prepare for remote snapshots is to create the management group and volumes that will be remote *before* taking the snapshot. Although the interface allows you to create volumes and snapshots as you go, that may be a distraction at the time a crucial snapshot is needed.

#### Getting there

This procedure takes you to the New Remote Snapshot window where remote copy procedures start.

1. In the navigation window, log in to the management group that contains the primary volume or snapshot for which you are creating the remote snapshot.

   You can create remote volumes and snapshots within the same management group. In that case, log in to the required management group.
2. Log in to the *remote* management group.
3. In the navigation window, select the primary volume (or snapshot).

   If you want to copy an existing snapshot to a remote management group, select that snapshot at this step.
4. Click **Snapshot Tasks**, and then select **New Remote Snapshot**.

## Creating the primary snapshot

1. In the Primary Snapshot Setup box, click **New Snapshot**.

   If you selected a snapshot to start the process, you do not need to create a new snapshot.
2. If you want to use VSS to quiesce the application before creating the snapshot, select the Application-Managed Snapshot check box.

   This option requires the use of the VSS Provider. For more information, see the *HP StorageWorks P4000 SAN Solution user guide* or online help. If the VSS Provider is not installed, use SAN/iQ to create a point-in-time snapshot (not using VSS).

   If the volume you selected is part of a volume set, see "Creating primary snapshots for volume sets" on page 19.

   This option quiesces VSS-aware applications on the server before SAN/iQ creates the snapshot.

   The system fills in the Description field and disables the Servers field automatically. You can assign servers after the snapshot is created.
3. Enter a name for the snapshot or accept the default.

   NOTE:

   Make the beginning of volume and snapshot names meaningful, for example, "Snap1Exchg_03."
4. (Optional) Enter a description of the snapshot.
5. Click **OK** to return to the New Remote Snapshot window.

   The information for the primary snapshot is pre-filled. For example, the text for the field Snapshot Name has changed from "Create Primary Snapshot" to "HdqtrsLogs_SS_1".
6. In the Remote Snapshot Setup box, use the drop-down lists to select the remote management group and volume.

   If you need to create a new volume, click **New Remote Volume**.
7. In the Snapshot Name field, enter the name for the remote snapshot.
8. (Optional) Enter a description for the remote snapshot.

9. Click **OK** in the New Remote Snapshot window.

   The remote copy of the primary snapshot to the remote volume begins.

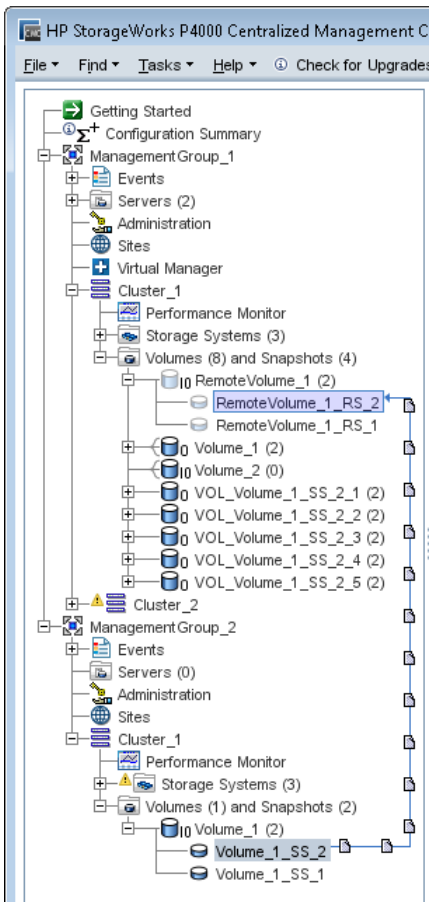


**Figure 4 Remote copy in progress**

## Creating primary snapshots for volume sets

The primary snapshot creation process for application-managed snapshots differs only when an application has associated volumes. Associated volumes are two or more volumes used by an application (volume set).

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.

When you create an application-managed snapshot of a volume in a volume set, the CMC recognizes that the volume is part of a volume set. The CMC then prompts you to create a snapshot for each volume in the volume set. This creates a snapshot set that corresponds to the volume set. To see any associated snapshots, select a snapshot, click the **Details tab**, and then look at the Snapshot Set field.

For information about the requirements for application-managed snapshots, see the *HP StorageWorks P4000 SAN Solution user guide* or online help.

NOTE:

After you create snapshots for a volume set, you do not want to delete individual snapshots from the snapshot set. You want to keep or delete all snapshots for the volume set. If you need to roll back to a snapshot, roll back each volume in the volume set to its corresponding snapshot. The system gives you the option to automatically delete or roll back all associated volumes.

To create a primary snapshots for volume sets:

1. Select a volume that is part of a volume set for the snapshot.
2. Log in to the management group that contains the volume for which you want to create a new snapshot.
3. Right-click on the volume and then select **New Remote Snapshot**.
4. Click **New Snapshot**.
5. Select the **Application-Managed Snapshot** check box.

   This option requires the use of the VSS Provider. For information about the requirements for application-managed snapshots, see the *HP StorageWorks P4000 SAN Solution user guide* or online help.

   This option quiesces VSS-aware applications on the server before SAN/iQ creates the snapshot.

   The system fills in the Description field and disables the Servers field automatically. You can assign servers after the snapshot is created.
6. Enter a name for the snapshot or accept the default.
7. Click **OK**.

   The New Snapshot—Associated Volumes window opens with a list of all volumes in the volume set.
8. (Optional) Edit the Snapshot Name for each snapshot.

NOTE:

Be sure to leave the Application-Managed Snapshots check box selected. This option quiesces the application before creating the snapshots. If you deselect the option, the system creates a point-in-time snapshot of each volume listed.

9. (Optional) Edit the Description for each snapshot.
10. Click **Create Snapshots** to create a snapshot of each volume.

    The Primary—Remote window opens with the list of the snapshots you just created. The first snapshot is already selected.
11. In the Remote Snapshot Setup box, use the drop-down lists to select the remote management group and volume for the selected snapshot.

NOTE:

If you need to create a new volume, select **New Remote Volume**.

12. In the Snapshot Name field, enter the name for this remote snapshot or accept the default.

13. (Optional) Enter a description for this remote snapshot.
14. Click **Update Pending Table Below** to add this remote snapshot setup to the list at the bottom of the window.

    The system selects the next volume in the Select Primary Snapshot list at the top of the window.
15. Select or enter the Volume Name, Snapshot Name, and Snapshot Description for this remote snapshot.

IMPORTANT:

All remote snapshots must be set up to use the same remote management group.

16. Click **Update Pending Table Below** to add this remote snapshot setup to the list at the bottom of the window.
17. Continue until each snapshot at the top of the window is set up.

    A green check mark shows the snapshot is set up.
18. Click **Create Remote Copies**.

    The remote copy of the primary snapshots to the remote volumes begins.

## Creating a remote volume

You can create a remote volume by using the following methods:

- Designate an existing primary volume as a remote volume.
- Create a new remote volume manually.
- Create a new remote volume during creation of a remote snapshot.
- Use the Management Groups, Clusters, and Volumes wizard in the "Getting Started" launch pad. See the *HP StorageWorks P4000 SAN Solution user guide* for details on working through the wizards.

### Designating an existing volume as a remote volume

When you select an existing volume to become a remote volume, the following occurs:

- A snapshot is created of the volume to preserve the data from the primary volume.
- The volume becomes a 0-byte remote volume.

### Creating a new remote volume manually

Create a remote volume as you would any other volume. Be sure to select the storage systems at the remote site. Because management groups and clusters are logical entities, name them to reflect their remote functionality.

In this method, the primary volume is ready. Create a remote volume at the remote site to receive the snapshot, and then either take the snapshot and create a remote copy, or create the schedule to take remote snapshots.

### Creating a remote volume while creating a remote snapshot

If you are using the New Remote Snapshot window, you can create a needed cluster and volume as you work through the window.

1. In the Remote Snapshot Setup box, select a Management Group to contain the remote snapshot.

NOTE:

You must be logged in to the management group you select.

2. Click **New Remote Volume**.

   The Management Groups, Clusters, and Volumes wizard opens.

   For specific help, see the "Getting Started" chapter in the *HP StorageWorks P4000 SAN Solution user guide* for details on working through the wizards.

   The wizard uses the information you entered to fill The New Remote Snapshot window when you exit the wizard.

3. (Optional) Enter a description of the remote snapshot and click **OK**.

   The system creates remote copy.

NOTE:

There may be a delay in remote copy time.

## What the system does

The system creates the remote snapshot in the cluster that contains the remote volume, and then copies the primary snapshot onto the remote snapshot. The process of copying the data may take some time.

The remote snapshot appears below the remote volume in the navigation window when the copy completes.

NOTE:

If you create a remote snapshot of a volume while a remote snapshot is in progress, the second remote snapshot does not begin copying until the first remote snapshot is complete.

## Creating the first copy

Creating the first copy of data is the first step when setting up a Remote Copy solution. Three methods for creating the first copy are described below.

**Copy data directly to the remote site over the WAN.**

Use this method if you are implementing the Remote Copy solution before you accumulate much data in the primary site, and your hardware is already installed in the remote site.

With this method, you create the primary management group and the remote management group in their respective locations, and then create the initial copy of the data directly over the WAN using Remote Copy.

**Use the storage systems intended for the remote site to configure the remote management group onsite and copy data locally, and then ship the remote storage systems to the remote site.**

Use this method if you initially have all the storage systems for the Remote Copy solution at the primary site.

1. Configure both the primary and remote management groups.
2. Create the first copy of the data locally over the gigabit Ethernet.
3. Ship the storage systems for the remote site and install the remote management group just as you configured it in the primary site.
4. Allow adequate time between the arrival of the storage systems and the first remote copy operation.

NOTE:

The subsequent snapshots from the primary volume to the remote volume are incremental.

**Use the PrimeSync method of Remote Copy to configure a temporary management group, create the first copy locally, ship the temporary storage system, and then copy locally to the remote target.**

Use this method if you have the primary (Site A) and remote site (Site B) configured and operational.

1. While at the primary Site A, use available storage systems to create a new temporary management group, cluster, and volume.

   This management group, cluster, and volume are the PrimeSync that you will set up at the primary Site A.
2. Make a remote snapshot of the primary Site A volume, and then copy it to the temporary PrimeSync management group over gigabit Ethernet. See "Creating a remote snapshot" on page 17.
3. Ship the storage systems to the remote Site B.
4. Power them on and discover them in the CMC to display the temporary PrimeSync management group, cluster, and volume.
5. Copy the remote snapshot from the temporary PrimeSync management group to the existing remote Site B management group.
6. Disassociate the temporary PrimeSync management group from the remote Site B management group. For more information, see "Disassociating remote management groups" on page 40.
7. Delete the temporary PrimeSync management group.
8. Set up the desired Remote Copy relationship, such as configuring a schedule to create remote snapshots of the volume from the primary Site A to remote Site B management group.

   PrimeSync ensures that the proper relationship is established between the original primary volume and the remote site. Subsequent remote snapshots from the primary site to the remote site are incremental.

NOTE:

Use the initial snapshot that you used for the temporary PrimeSync management group copy, to create the second Remote Copy or the schedule to create remote snapshots of the volume. You are now setting up the Remote Copy that goes from primary site A directly to remote site B which maintains that relationship going forward.

For more information on PrimeSync, see the "Application Note: SAN/iQ Remote Copy PrimeSync — Creating Initial Copy" at http://www.hp.com/support/manuals.

## Viewing a list of remote snapshots

View a list of remote snapshots associated with management groups, clusters, volumes, or snapshots.

1. In the navigation window, select a cluster to view its list of remote snapshots.
2. Click the **Remote Snapshot**.

   The report in the Remote Snapshot tab lists management groups and all the snapshots. The other columns show status information about the remote snapshots. For more information, see "Monitoring remote snapshots" on page 26.

## Setting the remote bandwidth

The remote bandwidth sets the maximum rate for data transfer between management groups. The copy rate is equal to, or less than, the rate set.

To control the maximum rate of data transfer to a remote snapshot, set the remote bandwidth on the management group that contains the remote snapshot which is the remote management group. When setting the remote bandwidth, you can choose from a list of common network types, or calculate a custom rate, based on your particular requirements.

### Selecting remote bandwidth rate

You may either select a preset speed from a list of standard network types or calculate a custom speed based on your specific requirements. Remember that the speed is the maximum rate at which data copies.

**Defaults setting**

When setting remote bandwidth, selecting Defaults allows you to choose from a list of common network types.

**Custom setting**

The custom setting for remote bandwidth defaults to 32,768 KB, or about 4 MB. Use the calculation tool to identify a desired bandwidth setting. For example, if you have a T1 line and you want to set the remote bandwidth to 12% of that capacity, you can use the calculation tool to find the correct value, 189 KB.

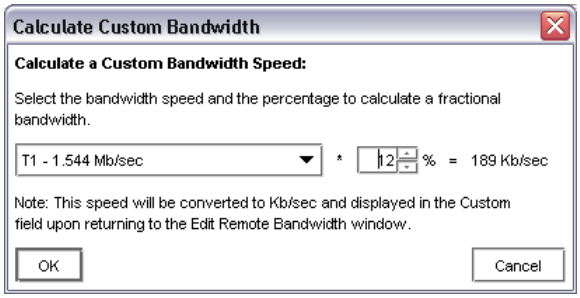


**Figure 5 Calculating a custom value for setting remote bandwidth**

### Best practice

Set the bandwidth speed the same in both directions unless you have an asymmetrical WAN link.

### Setting the bandwidth

1. In the navigation window, select either the remote or primary management group.
2. Click **Management Group Tasks**, and then select **Edit Management Group**.
3. Select the remote or primary management group.
4. Click **Edit Remote Bandwidth**.
5. Change the bandwidth setting as desired.

## Canceling a remote snapshot

When you cancel a remote snapshot that is in progress, the remote snapshot is deleted, but the primary snapshot remains.

1. In the navigation window, select the remote snapshot.
2. Click the **Remote Snapshot** tab.
3. Select the remote snapshot you want to cancel from the list, if it is not already selected.
4. Click **Remote Snapshot Tasks**, and then select **Cancel Remote Snapshot**.
5. Click **OK**.

## Editing a remote snapshot

To change the description and change the server assignment of a remote snapshot:

1. Log in to the management group that contains the remote snapshot.
2. Select the remote snapshot in the navigation window.
3. Click **Snapshots Tasks**, and then select **Edit Snapshot**.
4. Change the desired information, and then click **OK**.

## Deleting a remote snapshot

△ CAUTION:

Do not delete individual snapshots that are part of a snapshot set. To see associated snapshots, select a snapshot, click the **Details** tab, and then view the Snapshot Set field. For information about snapshot sets, see the *HP StorageWorks P4000 SAN Solution user guide* or online help. HP recommends that you keep or delete all snapshots for a volume set. If you need to roll back to a snapshot, you want to roll back each volume in the volume set to its corresponding snapshot. The system gives you the option to automatically delete or roll back all associated volumes.

1. Log in to the management group that contains the remote snapshot.
2. Select the remote snapshot in the navigation window.
3. Click **Snapshot Tasks**, and then select **Delete Snapshot** from the menu.

4. Click **OK**.

| If the snapshot is *not* part of a snapshot set | If the snapshot is part of a snapshot set |
|---|---|
| A confirmation message opens.<br>• Click **Delete Snapshot**. | A warning message opens.<br>• To delete all snapshots in the snapshot set, click **Delete All Associated Snapshots**.<br>• To delete only the snapshot you selected, click **Delete Selected Snapshot Only**.<br>• To cancel the deletion, click **Cancel**. |

# Monitoring remote snapshots

Information for monitoring remote snapshots is available from multiple sources. Remote snapshot alarms and events are listed in the Alarms window and in the Events category for the management group. You can also receive event notification through email and SNMP traps. For information about configuring event notification, see the *HP StorageWorks P4000 SAN Solution user guide*.

## Monitoring details from the Remote Snapshots feature

You can view information about each remote snapshot in both the Remote Snapshots tab and in the Remote Snapshot Details window.

### Viewing information in the Remote Snapshots tab

The Remote Snapshots tab displays a list of remote snapshots for each selected item in the navigation window. You can view lists of remote snapshots by management group, cluster, volume, and snapshot levels.

1. Select the appropriate item in the navigation window.
2. Click the **Remote Snapshot** tab.

The following fields appear:

- % Complete—The incremental progress of the remote copy operation
- Elapsed Time—The Incremental time of the copy operation
- Data Copied—The Incremental quantity of data copied
- Rate—The Rate at which data is being copied, or, when the remote snapshot is complete, the average rate for the total operation
- State—The Status of the operation

### Viewing the status in the Remote Snapshot Details window

The Remote Snapshot Details window displays additional details about a remote snapshot.

1. In the tab window, select the **Remote Snapshots** tab.
2. Select a remote snapshot from the list of snapshots on the Remote Snapshots tab.
3. Click **Remote Snapshot Tasks**, and then select **View Remote Snapshot Details**.

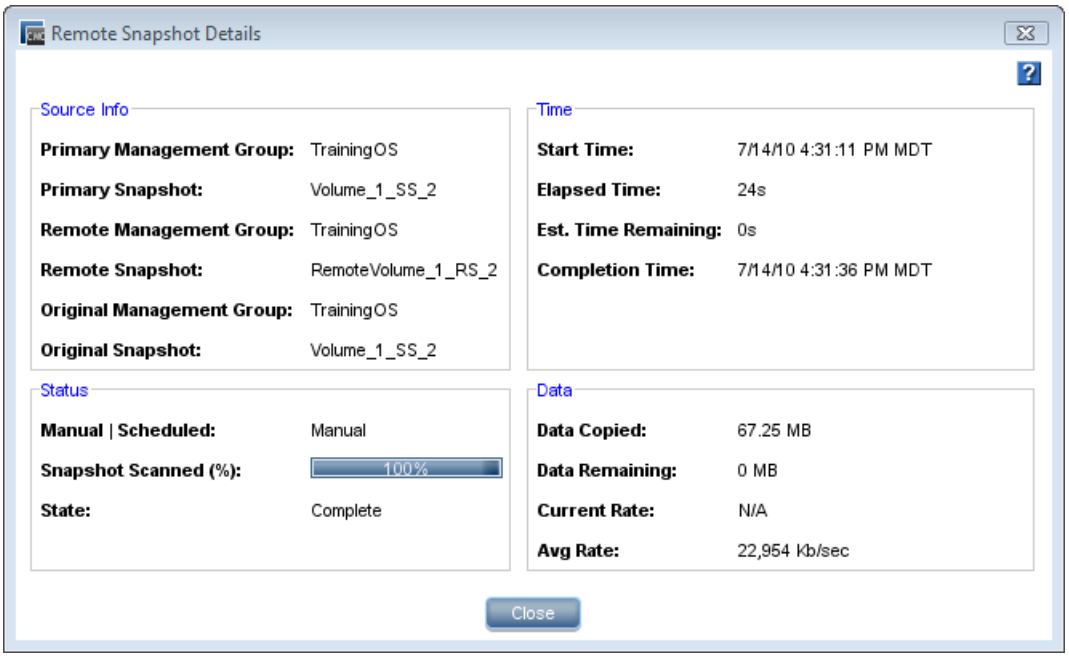


**Figure 6 Viewing remote snapshot details**

During the remote copy process, the Remote Snapshot Details window reports current statistics. When the copy is completed, the statistics show summary data. For field descriptions, see Table 5 on page 27.

**Table 5 Fields for Remote Snapshot Details window**

| Statistic | Fields |
|---|---|
| **Source Info section** | |
| Primary Mgmt Group | The management group containing the primary volume and snapshot |
| Primary Snapshot | The primary snapshot |
| Remote Mgmt Group | The management group containing the remote volume and snapshot |
| Remote Snapshot | The remote snapshot |
| Original Mgmt Group | The original management group that contained the original volume and snapshot. This is used with PrimeSync feature |
| Original Snapshot | The first version of the snapshot from which the first copy was created. This is used with PrimeSync feature. |
| **Status** | |
| Manual \| Scheduled | Whether the snapshot was created manually or with a scheduled snapshot |
| Snapshot Scanned (%) | The percentage complete of the copy process. Values are 0 to 100%. |
| State | The current state of the copy process. Valid values are Started, Copying, Stalled, and Complete. |
| **Time** | |

| Statistic | Fields |
|---|---|
| Start Time | Time Zone date and time copy started. The field format is MM/DD/YYYY HH:MM:SS [AM/PM] |
| Elapsed Time | The number of days, hours, minutes, and seconds since the copy began processing. The field format is Xd Xh Xm Xs X. N/A signifies completed copies or in-progress copies not yet calculated. |
| Est. Time Remaining | The number of days, hours, minutes, and seconds estimated to remain in the copy process. The field format is Xd Xh Xm Xs X. N/A signifies completed copies or in-progress copies not yet calculated. |
| Completion Time | Time Zone date and time copy completed. The field format is MM/DD/YYYY HH:MM:SS [AM/PM]. N/A signifies completed copies or in-progress copies not yet calculated. |
| **Data** | |
| Data Copied | Amount of data copied so far in smallest unit size |
| Data Remaining | Amount of data remaining to be copied in smallest unit size |
| Current Rate | Current rate of data being copied in Kb/second. This rate is recalculated regularly throughout the remote copy process. N/A if not yet available or completed. |
| Avg. Rate | Kb/sec. Average rate of copy progress. |

To monitor the progress of the remote copy, leave the Remote Snapshot Details window open.

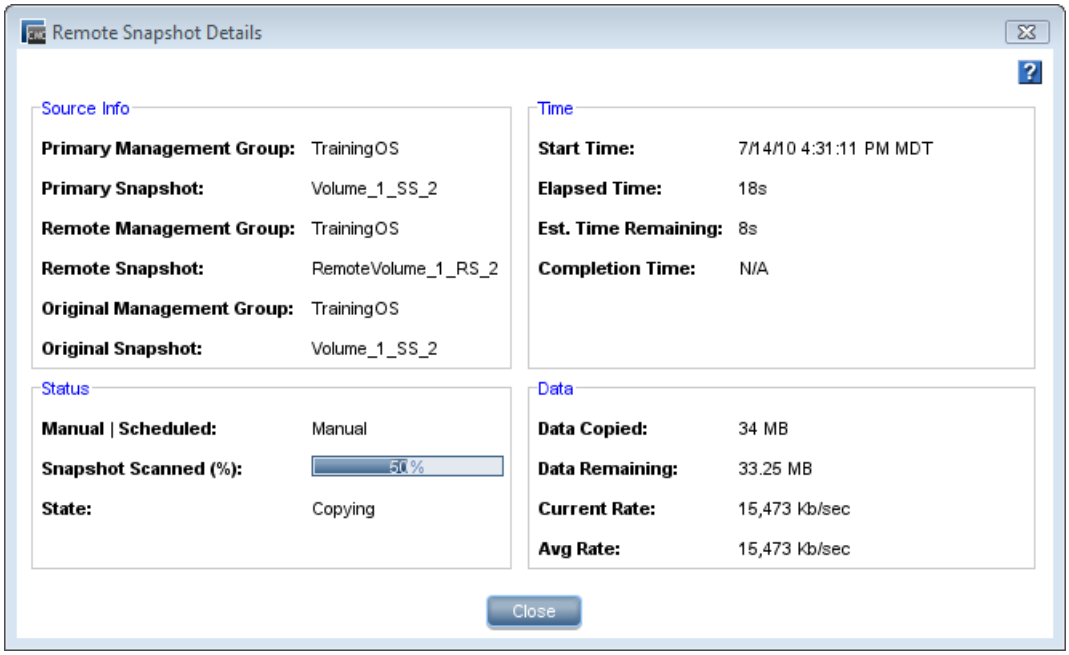


**Figure 7 Viewing remote snapshot details for remote copy in progress**

# Scheduling remote snapshots

In addition to taking remote snapshots of a volume manually, you can set up a schedule to take snapshots and save them remotely. Scheduled remote snapshots provide business continuance and disaster recovery, as well as a consistent, predictable update of data for remote backup and recovery.

Planning for scheduled remote snapshots is a crucial initial step in implementing Remote Copy. The following items require planning in advance for successful deployment of scheduled remote snapshots:

- Recurrence (frequency)
- Retention policies
- Capacity planning
- Timing

For detailed information about these issues, see "Planning for Remote Copy" on page 12.

# Best practices for scheduling remote snapshots

- Create a new remote volume to use with the scheduled remote snapshots.
- If you are performing daily remote copies, schedule the remote snapshots during off-peak hours. If you are setting scheduled remote snapshots for multiple volumes, stagger the schedules with at least an hour between start times.
- Use NTP to set all storage systems in the management group to the same time zone.
- Reset the management group time before creating a new schedule to remote snapshot a volume.

# Creating the schedule

To create the schedule for continuing remote snapshots:

1. In the navigation window, select the primary volume.
2. Click the **Schedules** tab.
3. Click **Schedule Tasks**, and then select **New Schedule to Remote Snapshot a Volume**.
4. Click **Edit**, and then select a **Start At** time.
5. Select a recurrence interval.

   You can schedule a snapshot to occur every 30 minutes or more.
6. Select a retention interval for the primary snapshot, either number of days or number of snapshots.

   You can retain up to 50 snapshots for a volume.
7. If you want to use VSS to quiesce the application before creating the snapshot, select the Application-Managed Snapshot check box.

   This option requires the use of the VSS Provider. For more information, see "Requirements for application-managed snapshots" in the *HP StorageWorks P4000 SAN Solution user guide*. If the VSS Provider is not installed, the SAN/iQ software allows the creation of a point-in-time snapshot.

   This option quiesces VSS-aware applications on the server before the SAN/iQ software creates the snapshot.
8. Select the management group and volume that will hold the remote snapshots.
9. Log in if you need to.

10. Click **New Remote Volume** to use the wizard to create a volume if you need to make a new one.
11. Set the retention interval for the remote snapshot.

    You can retain up to 50 snapshots for a volume.
12. Click **OK** to close the scheduling window and return to the navigation and tab windows.

    The timetable you just created is now listed in the Schedules tab view.

## Creating the schedule for volume sets

The schedule creation process for application-managed, remote snapshots differs only when an application has associated volumes. Associated volumes are two or more volumes used by an application (volume set).

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.

When you create an application-managed, remote snapshot of a volume in a volume set, the CMC recognizes that the volume is part of a volume set. SAN/iQ then prompts you to create a snapshot and remote copy of each volume in the volume set. This creates a snapshot set and remote copy that corresponds to the volume set. To see any associated snapshots, select a snapshot, click the **Details** tab, and then view the Snapshot Set field.

For information about the requirements for application-managed snapshots, see the *HP StorageWorks P4000 SAN Solution user guide* or online help.

When you first create the schedule, the system stores information about the volume set as it exists at that time. If you remove volumes from the volume set using the application, you must update the schedule. To update it, you only need to edit the schedule and click **OK**. The system automatically updates the volume set information when you click **OK**. If you want to see the updated information, you can click **Verify Volume Associations**, then click **OK**.

NOTE:

If you have a schedule to remote snapshot a volume and you add a volume to the volume set using the application, the system cannot update the volume set information as described above. You must delete the schedule and create a new one to reflect the current volume set.

The schedule also reflects the volume associations based on the volume you select when you create the schedule. That volume becomes the "owning" volume. The Volume Set field of the schedule displays (O) next to the owning volume. You should check that the field displays all of the volumes that you want to snapshot. It is possible that the owning volume is not aware of all associated volumes. If it is not, select a volume that is aware of all associated volumes and create the schedule there.

The procedure below assumes that you selected a volume that is part of a volume set.

1. In the navigation window, select the primary volume.
2. Click the **Schedules** tab.
3. Click **Schedule Tasks**, and then select **New Schedule to Remote Snapshot a Volume**.
4. Click **Edit**, and then select a **Start At** time.
5. Select a recurrence interval.

   You can schedule a snapshot to occur every 30 minutes or more.

6. Select a retention interval for the primary snapshot, either number of days or number of snapshots.

   You can retain up to 50 snapshots for a volume.

7. If you want to use VSS to quiesce the application before creating the snapshot, select the Application-Managed Snapshot check box.

   This option requires the use of the VSS Provider. For more information, see "Requirements for application-managed snapshots" in the *HP StorageWorks P4000 SAN Solution user guide*. If the VSS Provider is not installed, SAN/iQ allows creation of a point-in-time snapshot (not using VSS).

   This option quiesces VSS-aware applications on the server before SAN/iQ creates the snapshot.

8. Select the management group and volume that will hold the remote snapshots.
9. Log in if you need to.
10. Click **New Remote Volume** to use the wizard to create a volume if you need to make a new one.
11. Set the retention interval for the remote snapshot.

    You can retain up to 50 snapshots for a volume.

12. Click **OK**.

    The Volume Associations Found window opens with a list of all volumes in the volume set.

13. Review the information about the volume set and the remote copies the schedule will create.
14. Click **Continue**.

    The Primary and Remote Volume Associations window opens with the list of the primary volumes in the volume set. The first volume is already set up and the next one in the list is selected.

15. In the Remote Volume Setup box, use the Volume Name drop-down list to select the remote volume to use for the selected volume.

    All remote volumes must be set up to use the same remote management group. If you need to create a remote volume, click **New Remote Volume**.

16. Click **Update Pending Table Below** to add this remote volume setup to the list at the bottom of the window.

    The system selects the next volume in the Primary Volume list at the top of the window.

17. Select the remote Volume Name for the selected volume.
18. Click **Update Pending Table Below** to add this remote volume setup to the list at the bottom of the window.
19. Continue until each volume at the top of the window is set up.

    A green check mark shows the volume is set up.

20. Click **Create Schedule**.

    The timetable you just created is now listed in the Schedules tab view.

## Timing for a scheduled remote snapshot

When you set up a schedule for recurring remote snapshots with the previous procedure, check the time zone setting first. The time zone displayed in the Schedule to Remote Snapshot a Volume window is the time zone of the storage system through which you first logged in to the management group. See "Best practices for scheduling remote snapshots" on page 29.

## What the system does

If you created a new volume for the remote volume, the system creates a new primary snapshot of the primary volume and a remote snapshot of the remote volume. See "Best practices for scheduling remote snapshots" on page 29.

If you selected an existing volume to become the remote volume, the system alerts you that it will create a snapshot to save existing data on the volume. The snapshot that is then created retains all the volume's data.

1. Enter a name for that snapshot in the alert.
2. Click **OK** to continue.

   The new snapshot is created and the volume becomes a remote volume.

The system creates a new primary snapshot of the primary volume and a remote snapshot of the remote volume. It then copies the data from the primary snapshot to the remote snapshot. This process occurs according to the schedule.

# Pausing and resuming scheduled snapshots

At times, it may be convenient to prevent a scheduled snapshot from taking place. This section describes how to pause and then resume a schedule to snapshot a volume.

When you pause a snapshot schedule, the snapshot deletions for that schedule are paused as well. When you resume the schedule, both the snapshots and the snapshot deletions resume according to the schedule.

## Pause a schedule

1. In the navigation window, select the volume for which you want to pause the schedule.
2. Click the **Schedules** tab.
3. Select the schedule you want.
4. Click **Schedule Tasks** on the Details tab, and then select **Pause Schedule**.
5. In the Confirm window, click **OK**.

   In the Next Occurrence column of the Schedules tab window, this snapshot schedule is marked as paused.
6. Make a note to resume this snapshot schedule at a convenient time.

## Resume a schedule

1. In the navigation window, select the volume for which you want to resume the snapshot schedule.
2. Click the **Schedules** tab.
3. Select the schedule you want.
4. Click **Schedule Tasks** on the Details tab, and then select **Resume Schedule**.
5. In the Confirm window, click **OK**.

   In the Next Occurrence column of the tab window, this snapshot schedule shows the date and time the next snapshot will be created.

# Editing the schedule to remote snapshot a volume

When editing the timetable for a schedule to remote snapshot a volume, you can change the following items:

- Schedule—Description, start date and time, recurrence policy
- Primary Setup—Retention policy
- Remote Setup—Retention policy

## Edit the remote snapshot schedule

If the snapshot is part of a snapshot set, you can also verify that the volumes included in the schedule are the current volumes in the volume set. For more information, see "Understanding schedules to snapshot volume sets" in the *HP StorageWorks P4000 SAN Solution user guide*.

1. In the navigation window, select the primary volume that has the schedule you want to edit.
2. Click the **Schedules** tab, and then select the schedule to edit.
3. Click **Schedule Tasks**, and then select **Edit Schedule**.
4. Change the desired information.
5. (Optional) If displayed, click **Verify Volume Associations** to see if the volume set included in the snapshot set is up to date.

   The Verify Volume Associations window opens, and displays the volumes currently associated with the schedule. Any volumes that have been added to or removed from the volume set are reflected here.
6. Click **Close** to return to the Edit Schedule to a Snapshot Volume window.

   The updated list of volumes is populated in the Volume Set field. For more information, see "Understanding schedules to snapshot volume sets" in the *HP StorageWorks P4000 SAN Solution user guide*.

   The Volume Set field allows you to see the current volume set information.
7. Click **OK**.

   The information is automatically updated.

# Deleting the remote snapshot schedule

1. In the navigation window, select the primary volume that has the schedule you want to delete.
2. Click the **Schedule** tab.
3. Select the schedule you want to delete.
4. Click **Schedule Tasks**, and then select **Delete Schedule**.
5. Click **OK**.

# Failover and failback using Remote Copy

Configuring Remote Copy for failover provides for business continuance and disaster recovery. When configuring failover, consider both the failover and failback paths.

## Planning failover

To achieve failover, consider the following points:

- The location and structure of management groups and clusters
- Configuration of primary and remote volumes, snapshots, and scheduling snapshots
- Configuration of application servers and backup application servers
- Task flow for failback (resuming production after failover)
- If a volume is part of a volume set, typically you want to fail over each volume using its corresponding snapshot. To see associated snapshots, select a snapshot, click the **Details** tab, and review the Snapshot Set field. For more information about volume sets and snapshot sets, see the *HP StorageWorks P4000 SAN Solution user guide* or online help.

### Using scripting for failover

Application-based scripting provides the capability for creating, mounting, and deleting snapshots using scripts. Remote Copy can be scripted as well. Remote snapshots and scheduled remote snapshots can be created and managed using scripts. For information about scripting, see the *CLIQ—The SAN/iQ Command-Line Interface User Manual* and the sample scripts available at http://www.hp.com/go/P4000downloads.

# Using the Volume Failover/Failback Wizard

Changing the roles of primary and remote volumes may be necessary during failover and failback. Move your primary volume either as part of a failover/failback scenario or as part of a planned move. The Volume Failover/Failback Wizard takes you through the necessary steps to move a primary volume to an existing remote volume, and to make the existing remote volume to an acting primary volume.

Later, when failing back or restoring operations in a planned move, resynchronize data between the acting primary and the recovered, or newly configured, production site primary volume.

NOTE:

When failing over a primary volume, the volume must have a remote copy relationship for the wizard to be available.

## Converting a primary volume to a remote volume

The system allows you to convert a primary volume into a remote volume. First, the system takes a snapshot of the primary volume to preserve the existing data that is on the volume. Next, the primary volume is converted to a remote volume.

1. In the navigation view, select the volume that you want to convert.
2. Right-click the volume, and then select **Failover/Failback Volume**.
3. Click **Next**.
4. Select the reason you are failing over the volume.

   Use the first choice if your primary volume is available and you are planning a preemptive move of the primary volume. For the steps to use the second choice, see "Using failover to remote volume" on page 35.

5. Select the first choice, “To move the primary volume” and click **Next**.

   The next window reminds you to disconnect any iSCSI sessions connected to the volume.

6. Disconnect the iSCSI sessions, if any are displayed, and then click **Next**.

   The next window begins the process to make the primary volume into a remote volume.

7. Enter a name and a description for the snapshot to be taken of the current primary volume.

   This snapshot preserves the existing data on the volume.

8. Click **Next**.

9. The next window is where you designate the destination for copying the snapshot to a remote snapshot.

10. The final step is to make the remote volume into an acting primary volume.

    This acting primary volume connects to application servers in order to maintain business continuance or accomplish disaster recovery.

NOTE:

You cannot make a remote volume into a primary volume while a remote snapshot is in progress. Either wait until the remote snapshot copy completes before making the remote volume into a primary volume, or cancel the in-progress remote copy.

11. Click **Finish**.

    The snapshot is created and the volume becomes a remote volume.

The final window of the wizard displays a summary of the actions and a reminder to reconnect your iSCSI sessions.

## Using failover to remote volume

If the primary volume is not available, you can use the wizard to promote the remote volume to an acting primary volume.

1. In the navigation view, select the volume that you want to convert.
2. Right-click the volume, and then select **Failover/Failback Volume**.
3. Click **Next**.
4. Select the reason you are failing over the volume.

   Use the second choice if your primary volume is not available and you want to get an acting primary volume into production.

5. Select **To failover to the remote volume**, and click **Next**.

   The next window reminds you to disconnect any iSCSI sessions connected to the volume.

6. The final step is to make the remote volume into an acting primary volume.

   This acting primary volume connects to application servers in order to maintain business continuance or accomplish disaster recovery.

7. Click **Finish**.

   The final window of the wizard displays a summary of the actions and a reminder to reconnect your iSCSI sessions.

8. If you promoted a remote application-managed snapshot, use diskpart.exe to change the resulting volume's attributes.

   For more information, see the *HP StorageWorks P4000 SAN Solution user guide* or online help.

# Resuming production after failover

After failover occurs, three scenarios exist for resuming production:

- Failback returns operations to the original primary site it is restored.
- Convert the backup site to the new primary site.
- Set up a new primary site and resume operations at that site.

The task flow for restoring or recovering data and resuming the original Remote Copy configuration is different for each scenario.

Use these procedures when you are resynchronizing data between the acting primary volume and the recovered, or newly configured, production site primary volume.

## Synchronizing data after failover

After a failover, there are usually two snapshots or volumes that have conflicting data. Recovering and synchronizing data depends on multiple factors, including the application involved.

### Example scenario

The following example illustrates one process for synchronizing data. Remember that synchronization is optional.

**Table 6 Timeline of failover**

| Time | Event | What happens |
|---|---|---|
| 1:00 p.m. | Regular hourly scheduled remote snapshot starts. | `RemoteSS_1` creates in remote management group. |
| 1:10 p.m. | Remote copy finishes. | Copying is complete. |
| 1:30 p.m. | Primary volume goes offline. | `OrigPrimaryVol_1` is offline. |
| 1:33 p.m. | Scripted failover causes remote volume to become the acting primary volume. | `ActPrimaryVol_1` becomes primary and active, and usable to application server. |
| 2:00 p.m. | Original primary volume comes back online. | `OrigPrimaryVol_1` is online. |

The following data needs to be synchronized once the primary volume is back online:

- Original volume, which contains data from 1:00 to 1:30 p.m.
- Acting primary volume which contains data from 1:33 to 2:00 p.m.

## Returning operations to original primary site

After the original primary site is operational again, restore operations to that site. The steps to restore operations depend upon the state of the original primary volume.

- If the primary volume is working:

Synchronize the data between the acting primary volume and the restored primary volume before returning the acting primary volume to its remote volume state.

- If the primary volume is not available:

  Create a new primary volume, synchronize the data with the acting primary volume, and then return the acting primary volume to a remote volume.

## Synchronizing the data between the acting primary volume and the original primary volume

**Create snapshots that contain the data that you need to synchronize.**

**Table 7 Creating snapshots of data to synchronize**

| Action | Volumes and snapshots on primary management group | Volumes and snapshots on remote management group | What this step accomplishes |
|---|---|---|---|
| 1. Stop applications that are accessing the volumes. | | | |
| 2. Make a snapshot of the original volume. | `OrigPrimary-Vol_1OrigPrimarySS_1` | | Creates a snapshot of the original primary volume that includes the data from 1:00 - 1:30 p.m. |
| 3. Make the acting primary volume into the remote volume. This automatically creates a snapshot of the acting primary volume. | | `Remote-vol_1ActPrimarySS_1` | Returns the remote management group to its original configuration. In addition, you capture the 1:33 to 2:00 p.m. data. |

**Synchronize the data**

Synchronize the snapshots `OrigPrimarySS_1` and `ActPrimarySS_1` that were created in Steps 2 and 3 of Table 7 on page 37.

In the simplest case, to synchronize the snapshots, remote copy the remote snapshot back to the original primary volume.

## Creating a new primary volume at the original production site

If the original primary volume is not available, designate a new primary volume, synchronize the data from the acting primary volume, and then configure the timetable for the scheduled remote snapshot schedule on the new primary volume.

1. Stop the application that is accessing the acting primary volume.
2. Create a remote snapshot of the acting primary volume.
3. As your target, create a remote volume, which will later be converted into your primary volume.
4. Convert the remote volume into a primary volume.
5. Make the acting primary volume into the remote volume.

   This creates a snapshot of that volume.
6. Configure a new timetable for the scheduled remote snapshots on the new primary volume.

7. Reconfigure scripts for failover on the application servers.

## Setting up a new production site

Setting up a new production site involves creating a new primary volume and synchronizing the acting primary volume before returning it to its original state as a remote volume. The steps are the same as those for creating a new primary volume at the original production site.

## Making the backup site into the new production site

Turn the backup site into the new production site and designate a different backup site. The steps are similar to those for initially configuring Remote Copy.

1. Create a remote snapshot or a timetable for a scheduled remote snapshot on the acting primary volume.
2. Make a new remote volume on the new backup site as part of creating that remote snapshot or timetable for a scheduled remote snapshot.
3. Reconfigure scripts for failover on the application servers.

# Rolling back primary and remote volumes

Rolling back a volume from a snapshot is one method for reverting to an earlier copy of the data on a volume. Rolling back procedures requires that you delete any snapshots that were created after the snapshot that is rolled back to.

## Using SmartClone to clone a volume from a snapshot

Consider using the SmartClone feature to clone a volume from a snapshot that contains the earlier copy of the data you want to use. Creating a SmartClone volume preserves all snapshots while providing an exact copy of the desired data. The SmartClone volume consumes no extra space on the SAN.

### Rolling back a primary volume

Rolling back a primary volume to a snapshot replaces the original volume with a volume that contains the snapshot's data. The new volume has the same name as the original.

If a volume is part of a volume set, typically you want to roll back each volume using its corresponding snapshot. The system allows you to automatically roll back all associated volumes. To see any associated snapshots, select a snapshot, click the **Details** tab, and then look at the Snapshot Set field. For more information about volume sets and snapshot sets, see the *HP StorageWorks P4000 SAN Solution user guide* or online help.

### Prerequisites

Stop applications from accessing the volume.

△ CAUTION:

Any remote snapshot that has not completed copying is canceled.

1. Log in to the management group that contains the primary volume that you want to roll back.
2. Select the snapshot that you want to roll back to.
3. Review the snapshot Details tab to ensure you have selected the correct snapshot.
4. Click **Snapshot Tasks** on Details tab, and then select **Roll Back Volume**.

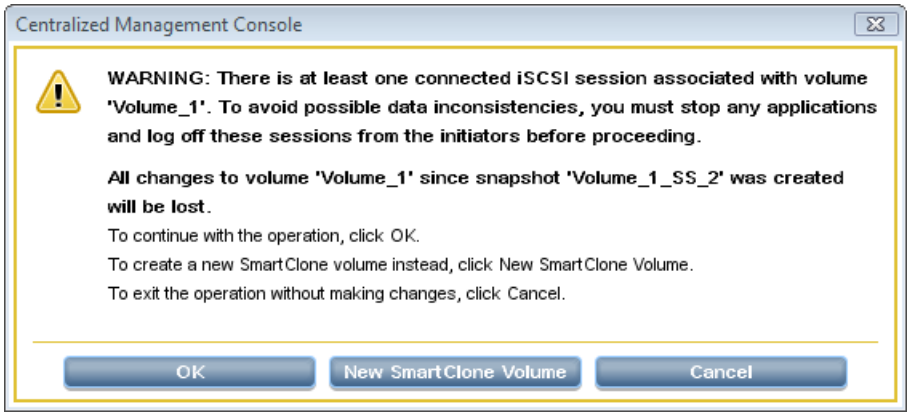


**Figure 8 Rolling back a primary volume if iSCSI sessions are still connected**

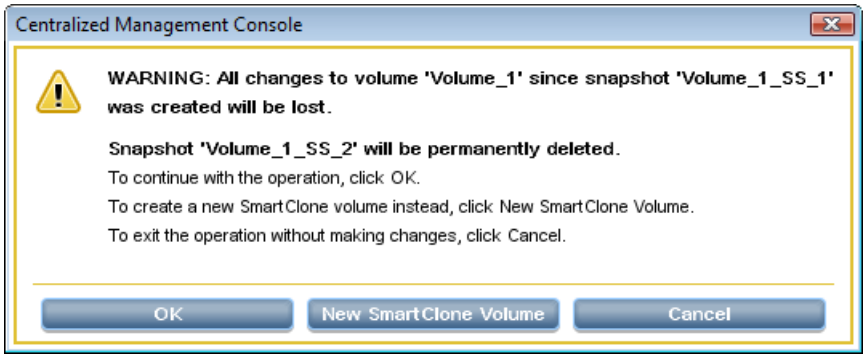


**Figure 9 Rolling back a primary volume if iSCSI sessions are not connected**

5. Log off any connected iSCSI sessions.
6. Click **OK**.

   The primary snapshot version of the primary volume is restored as a volume with the same name.
7. Reconfigure application servers to access the new volume.

## Rolling back a remote volume

A remote volume cannot be rolled back until you convert the remote volume into a primary volume and then perform the steps in Rolling back a primary volume.

If a volume is part of a volume set, typically you want to roll back each volume using its corresponding snapshot. The system gives you the option to automatically roll back all associated volumes. To see any associated snapshots, select a snapshot, click the **Details** tab, and look at the Snapshot Set field. For more information about volume sets and snapshot sets, see the *HP StorageWorks P4000 SAN Solution user guide* or online help.

# Using remote snapshots for data migration and data mining

Use remote snapshots to create split mirrors for data mining and data migration. A split mirror is a one-time remote snapshot created from the volume containing the data you want to use or move. Split mirrors are usually created for one-time use and then discarded.

## New for release 8.0

Consider using the SmartClone feature to create a split mirror. The SmartClone volume is an exact copy of the volume or snapshot yet consumes no extra space on the SAN.

### Creating a split mirror

To create a split mirror:

- Create a remote snapshot.
- Create a server for client access.
- Configure clients to access the remote snapshot.

# Disassociating remote management groups

Management groups become associated when linked by either remote snapshots or scheduled remote snapshots. Disassociating management groups destroys all the shared knowledge between those groups.

## Best practice for disassociating management groups

NOTE:

Do this only if a group no longer exists, or if instructed by Customer Support.

1. Log in to both management groups that you want to disassociate.
2. In the navigation window, select the remote management group.
3. Click **Management Group Tasks**, and then select **Edit Management Group**.
4. Select the management group or groups you want to disassociate, that is, the management groups that are remote relative to this management group.
5. Click **Disassociate**.

   A confirmation message opens, describing the results of disassociating the management groups.

CAUTION:

Disassociating the management group cancels any in-progress remote snapshots and deletes all timetables between the primary and remote management groups.

1. Click **OK**.

   The Edit Management Group window is displayed on top again, and the remote management group you disassociated from is gone from the list.
2. Click **OK** to return to the navigation window.

# 3 Sample Remote Copy configurations

Because of its flexibility, Remote Copy is useful in a variety of configurations. The sample configurations described in this chapter show only a few ways to use Remote Copy for business continuance, backup and recovery, data migration, and data mining.

## Using Remote Copy for business continuance

Business continuance is composed of disaster recovery and high availability of data. If using Remote Copy for business continuance, data is stored off-site and is readily available in the event of a site or system failure.

### Achieving high availability

Creating remote snapshots in remote locations with Remote Copy ensures that applications such as SQL Server, Oracle, and Exchange have access to backup copies of data volumes if production application servers or data volumes fail.

Using off-site remote snapshots of your production volumes, you can configure a backup application server to access those remote snapshots or volumes. Off-site remote snapshots, particularly when supplemented with synchronous Network RAID within a cluster, ensure high availability of critical data volumes.

### Configuration for high availability

To use remote snapshots for high availability, configure a backup application server to access remote volumes in the event of a primary system failure. Figure 10 on page 44 illustrates this simple high availability configuration.

- Configure primary and backup application servers.

  During normal operation, the production application server reads and writes to the primary volume.
- Set up a schedule for copying remote snapshots to the backup location. If your application server uses multiple volumes that must be in sync, use a script or VSS to quiesce the application before creating remote snapshots.

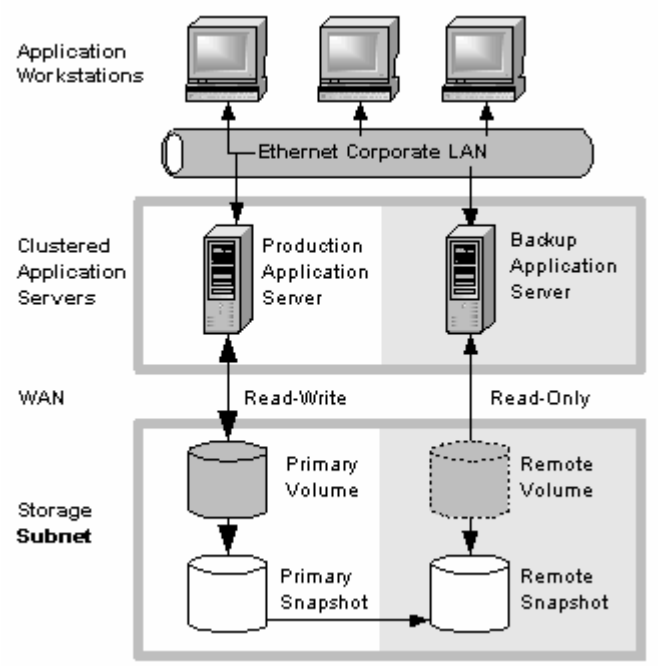


**Figure 10 High availability example configuration**

## How this configuration works for high availability

If the production application server or volumes become unavailable, application processing fails over to the backup application server. The remote volume and remote snapshots become acting primary, and the backup application server becomes the acting production application server, accessing data from the acting primary volume.

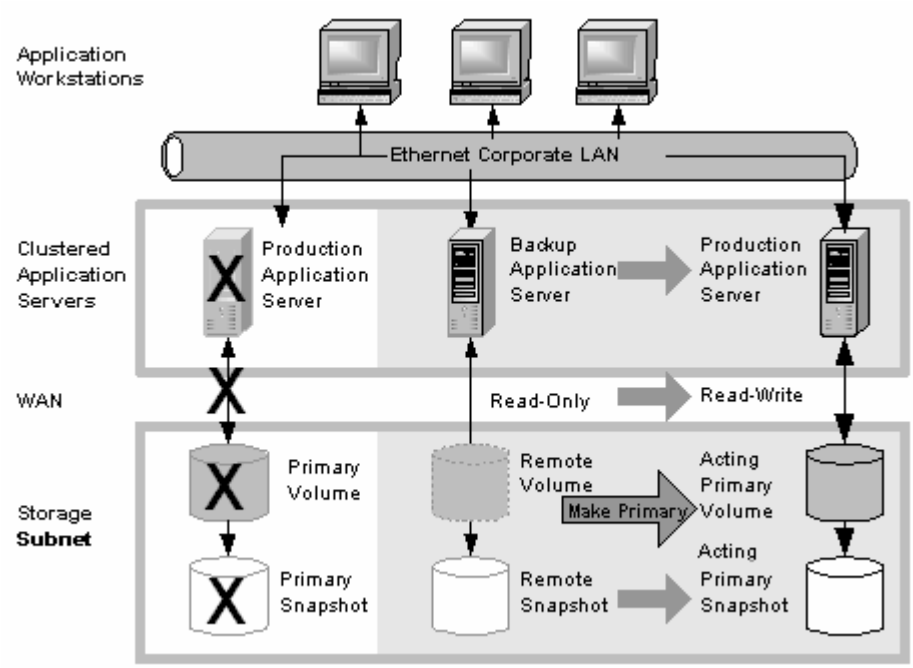


**Figure 11 High availability configuration during failover**

### Data availability if the primary volume or production application server fails

If either the primary volume or production application server in your production site fails, only that data written to the volume since the last remote snapshot was created is unavailable until the volume or production application server is restored.

### Failover to the backup application server

To maintain availability of the applications and the remaining data, the following process occurs:

1. A script or other application monitoring the production application server discovers that the primary volume is not available. A script executes to fail over to the backup application server.
2. The backup application server executes a script to convert the remote volume into a primary volume so that the volume can be accessed by the backup application server. For information about scripting, see the *CLIQ—The SAN/iQ Command-Line Interface User Manual* and the sample scripts available at http://www.hp.com/go/P4000downloads.
3. Because the backup application server was configured to access the remote (now primary) volume, operation of the backup application server begins.

The application continues to operate after the failover to the backup application server.

### Failback to the production configuration

When the production server and volumes become available again, you have the following two failback options:

- Resume operations using the original production server, and then return the backup volumes to their original remote status, as illustrated in Figure 12 on page 46. This requires migration of data that was written to the backup volumes since the failure back onto the production volumes.
- Continue operating on the backup application server. When the production server and volumes become available, configure the production server to be the backup server (role reversal).

### Merging data for failback

In the failover scenarios described above, there are probably two snapshots with different data. As part of failback, users must make a decision whether to merge the data from the two snapshots and the most effective method for doing so. For more information see "Synchronizing the data between the acting primary volume and the original primary volume" on page 37.

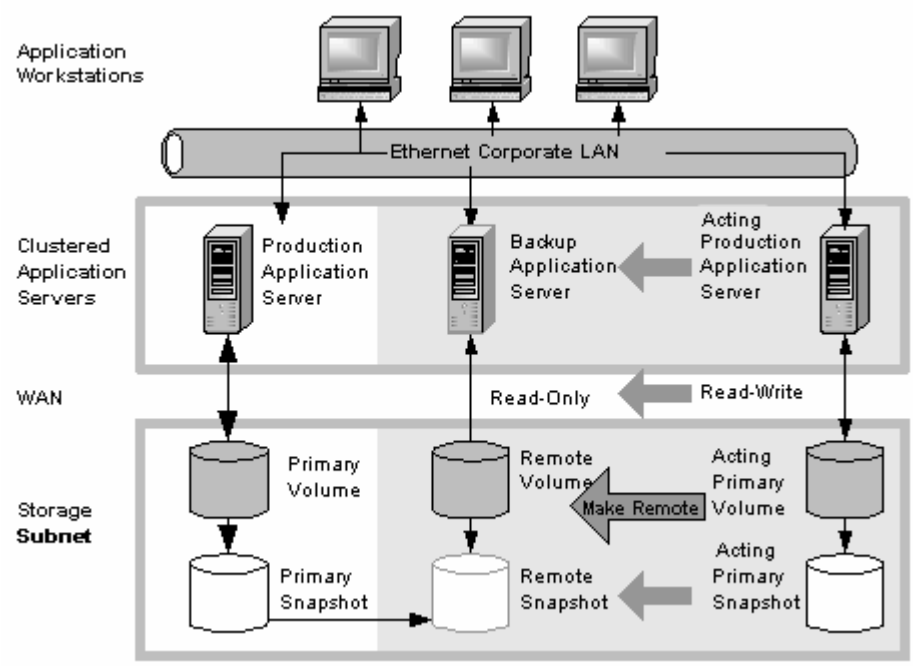


**Figure 12 High availability configuration during failback**

# Best practices

## Remote snapshots with volume replication

Use remote snapshots in conjunction with local, synchronous volume replication, known as Network RAID. Using remote snapshots alone, any data written to the primary volume since the most recent remote snapshot was created will be unavailable if the primary volume is unavailable.

However, you can lessen the impact of primary volume failure by using Network RAID. Network RAID allows you to create up to four copies of a volume on the same cluster of storage systems as the primary volume. The only limitation is that the cluster must contain at least as many storage systems as replicas of the volume. Replicating the volume within the cluster ensures that if a storage system in the cluster goes down, replicas of the volume elsewhere in the cluster will still be available. For more information about Network RAID and data protection levels, see the chapter "Provisioning Storage" in the *HP StorageWorks P4000 SAN Solution user guide*.

## Example configuration

Figure 10 on page 44 uses three storage systems per cluster. However, this scenario can use any number of storage systems. For information about creating clusters and volumes, see the *HP StorageWorks P4000 SAN Solution user guide*.

- In the production location, create a management group and a cluster of three storage systems.
- Create volumes on the cluster and set the data protection level to Network RAID-10.
- Configure the production application server to access the primary volume via iSCSI.
- In the backup location, create a second management group and a cluster of three storage systems.
- Create a schedule for making remote snapshots of the primary volume. See "Scheduling remote snapshots" on page 29.

NOTE:

Data protection levels are set independently for primary and remote volumes.

### How it works

If one of the storage systems in the primary location fails, the primary volume will still be available. If all of the storage systems fail, or if the application server fails, then failover to the backup application server occurs, and the remote snapshot(s) becomes available.

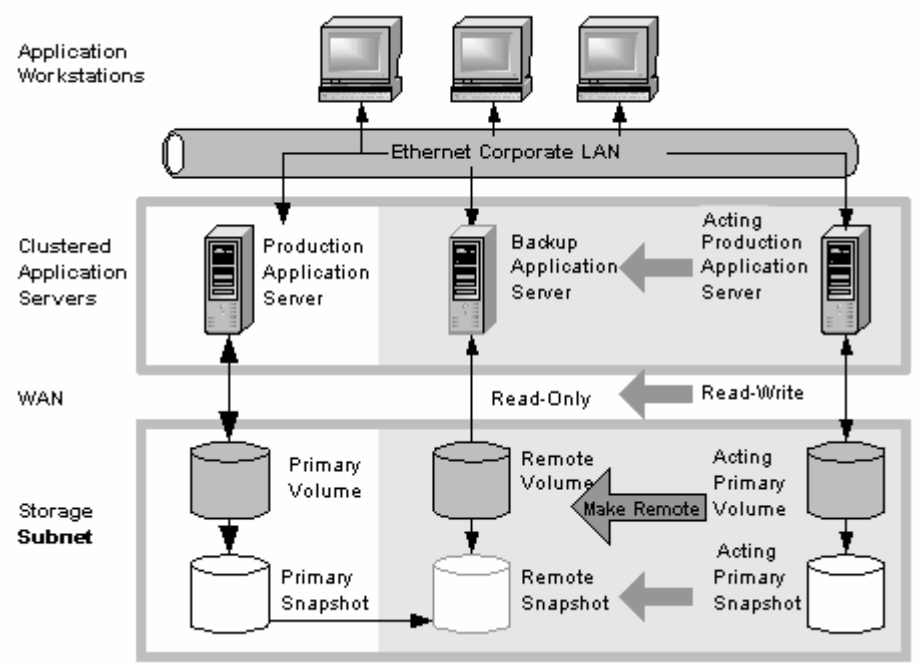


**Figure 13 High availability during failover-Example configuration**

## Achieving affordable disaster recovery

Even if you do not have clustered application servers or network bandwidth required for configuring hot backup sites, you can still use Remote Copy to protect your data during an emergency.

Using remote snapshots, you can maintain copies of your volumes in remote sites. Set up a schedule for creating remote snapshots, and if your primary storage site becomes unavailable, you can easily access the most recent remote copy of your data volumes. You can also use remote snapshots to transfer data to a backup location where tape backups are then created. This eliminates the backup window on your primary volumes, and ensures that you have copies of your data in the remote site on storage systems as well as on tape.

## Configuration for affordable disaster recovery

To configure affordable disaster recovery, create remote snapshots of your volumes in an off-site location. In addition, you can create tape backups from the remote snapshots in the off-site location:

- Designate one or more off-site locations to be the destination for remote snapshots.

- Set up a schedule for creating remote snapshots in the designated off-site locations. If your application server uses multiple volumes that must be in sync, use a script to quiesce the application before creating remote snapshots.
- Create routine tape backups of the remote snapshots in the off-site locations.

Figure 14 on page 48 shows an example configuration for disaster recovery.

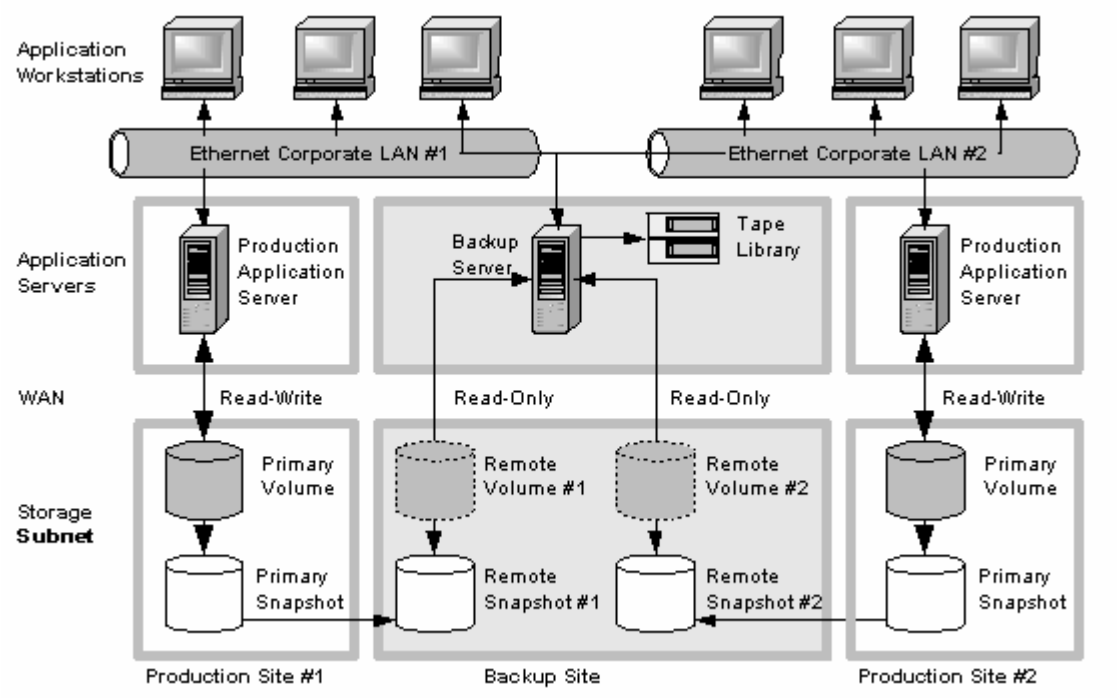


**Figure 14 Affordable disaster recovery example configuration**

## How this works for affordable disaster recovery

If the storage systems in your primary location fail or volumes become unavailable, the off-site location contains the most recent remote snapshots.

- Use the remote snapshots to resume operations as shown in Figure 15 on page 49. If you created tape backups, you can recover data from tape backups, as shown in Figure 16 on page 49.
- Only data written to the primary volumes since the last remote snapshot was created will be unavailable.
- Application servers that were accessing the offline volumes will not be available until you reconfigure them to access recovered data.

To resume operations using the most recent set of remote snapshots:

1. In the backup location, make the remote volume into a primary volume.

2. Configure application servers to access the primary volume, or if network connections are not fast enough to facilitate reading and writing to the off-site location, copy this volume to a location where application servers can access it more efficiently.

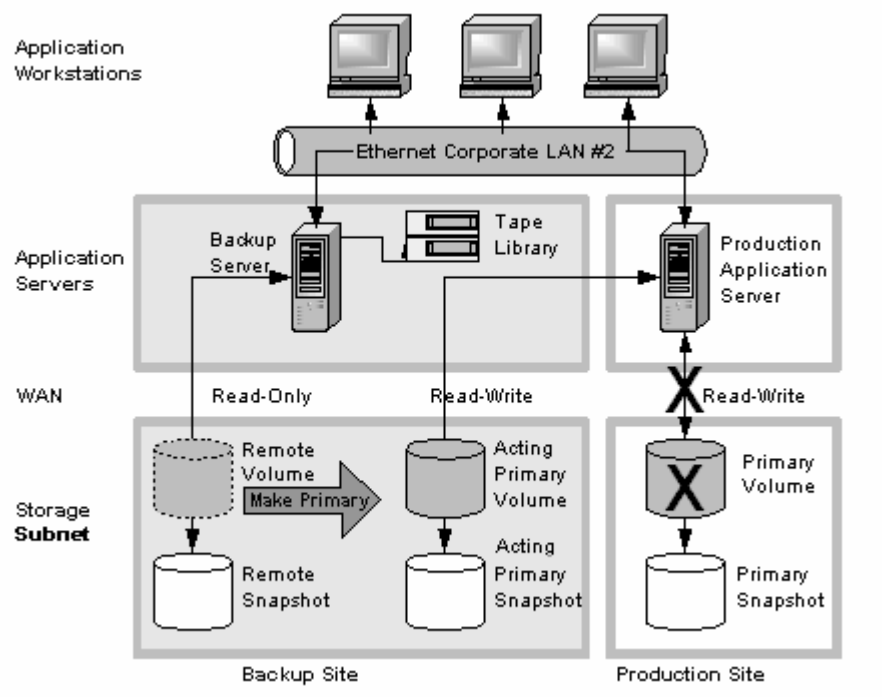


**Figure 15 Restoring from a remote volume**

In Figure 15 on page 49 note the volume labeled Primary Snapshot in the Production Site. It originated as a read only backup, but is brought into use as an acting primary.

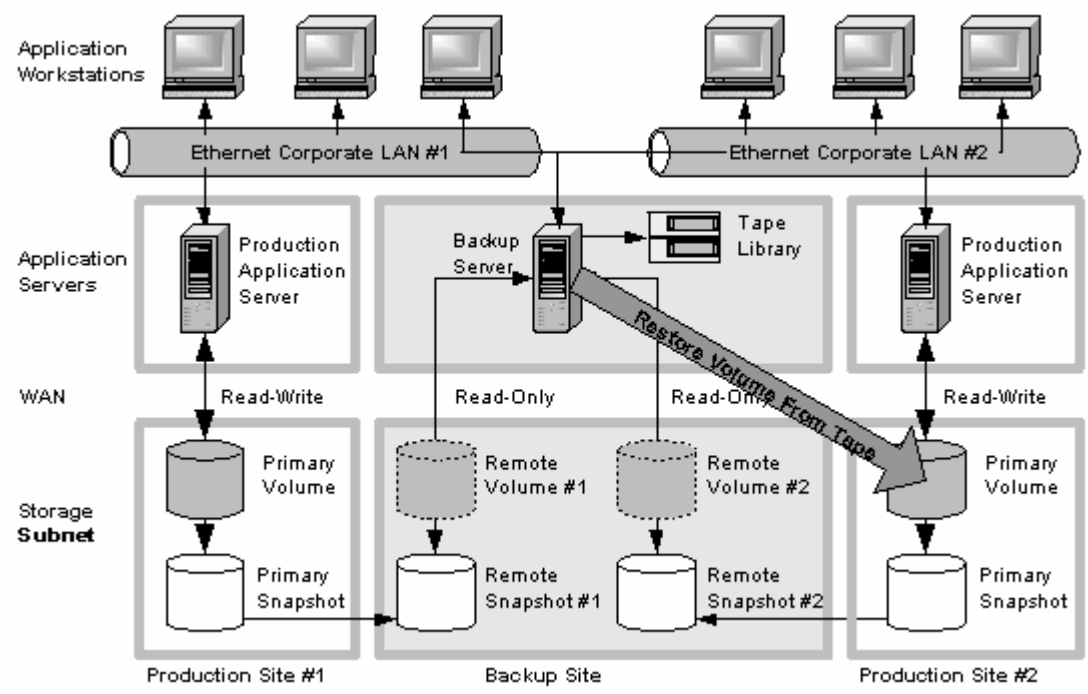


**Figure 16 Restoring from tape backup**

## Best practices

### Select an optimum recurrence schedule.

Select a recurrence schedule for remote snapshots that minimizes the potential for data loss. Any data written to the primary volume since the most recent remote snapshot was created will be unavailable if the primary volume is unavailable. Consider how much data you are willing to lose in the event of an emergency and set the recurrence for creating remote snapshots accordingly.

If you do not want a large number of remote snapshots to accumulate on your remote volume, you can use several timetables for scheduled remote snapshots, each with different retention policies. For example, suppose you want to create remote snapshots every four hours to ensure that no more than four hours worth of data is lost in an emergency. In addition, you want to retain one week's worth of remote snapshots. Retaining four-hour snapshots for one week can result in the accumulation of over 40 remote snapshots. Another approach would be to create two remote snapshot schedules for the volume:

- One schedule to create remote snapshots every four hours, but only retain the most recent six remote snapshots. This ensures that you do not lose more than four hours worth of data in an emergency.
- A second schedule to create remote snapshots every 24 hours and retain 7 remote snapshots.

### Use remote snapshots in conjunction with local, synchronous volume replication

To prevent data loss, reinforce Remote Copy with Network RAID within the cluster of storage systems at the primary geographic site. With Network RAID configured, a single storage system can be off-line, and your primary volume remains intact.

At the backup location, you can also use synchronous replication to protect your remote volume against storage system failure.

#### Example configuration

- In the production location, create a cluster of three storage systems, all with managers.
- Create volumes on the cluster, and set the data protection level to Network RAID-10.
- Create a schedule for making remote snapshots of the primary volume. Set the recurrence to every four hours, and retention of remote snapshots to two days.

NOTE:

You can use the same data protection levels on the remote volume as well. However, this data protection level is configured independently of the data protection level that is configured on the primary volume.

If one of the storage systems in the primary location fails, the primary volume will still be available. If all of the storage systems fail, or if the application server fails, then you can recover data from the remote snapshots or tape backups in the off-site location.

# Using Remote Copy for off-site backup and recovery

For backup and recovery systems, Remote Copy can eliminate the backup window on an application server. Using iSCSI command line interface commands and scripts, configure the iSCSI initiator to mount remote snapshots on a local or remote backup server, and then back up the remote snapshot from the backup server. The remote snapshot is available if the primary volume fails.

## Achieving off-site backup

Rather than creating tape backups and then transporting them to a secure off-site location, you can use Remote Copy to create remote snapshots in an off-site location. Then, optionally, you can create tape backups at the off-site location.

## Configuration for off-site backup and recovery

To use remote snapshots for off-site tape backup, create remote snapshots for access by your tape backup application:

- Create remote volumes in your backup location.
- Configure your backup application to access the remote snapshots.
- Configure schedules to create remote snapshots in the designated off-site locations. If your application server uses multiple volumes that must be in sync, use a script to quiesce the application before creating remote snapshots.
- [Optional] Create routine tape backups of the remote snapshots.

### Configuration diagram

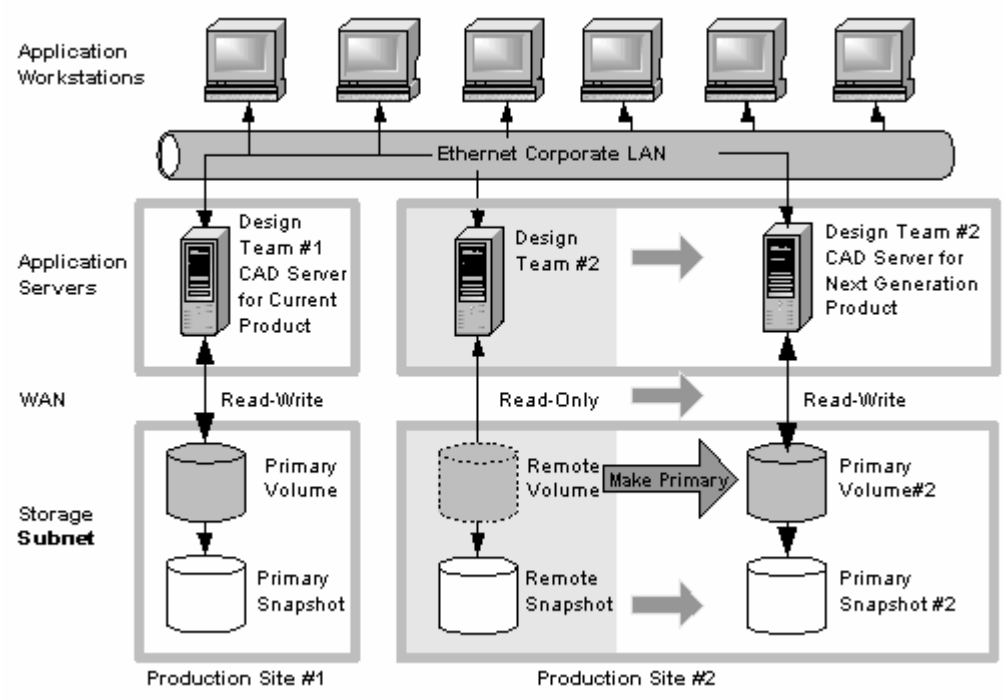


**Figure 17 Off-site backup and recovery example configuration**

## How this configuration works for off-site backup

Depending on how long you retain the copies of the remote snapshots, you can retrieve data directly from recent remote snapshots rather than using tape backups. Otherwise, retrieve data as you normally would from the tape backup.

## Best practices

**Retain the most recent primary snapshots in the primary cluster**

By keeping snapshots on your primary volume, you can quickly roll back a volume to a previous snapshot without accessing off-site backups.

- When you create a schedule for Remote Copy, specify the number of primary and remote snapshots that you want to retain. You can retain primary snapshots to facilitate easy rollback of the primary volume.

NOTE:

Retention of snapshots affects the amount of space that is used in the cluster of storage systems. Balance the number of snapshots to retain with the amount of space you are willing to use. You can still access remote snapshots or tape backups if you want to roll back to a snapshot that you did not retain.

- Retain remote snapshots in the backup location to facilitate fast recovery of backed-up data. If you retain a number of remote snapshots after a tape backup is created, you can access the data without going to the backup tape.

## Example configuration

- Retain three primary snapshots. This enables you to roll the primary volume back, yet it requires a relatively small amount of space on the primary cluster.
- Retain up to a week's worth of remote snapshots on the backup cluster.
- For snapshots older than one week, go to the backup tape.

# Achieving nondestructive rollback

As discussed in "Rolling back primary and remote volumes" on page 38, rolling a snapshot back to a volume requires you to delete any snapshots that were created since the snapshot that you roll back to. For example, suppose you created snapshots of a volume on Monday, Tuesday, and Wednesday. On Thursday, if you roll the volume back to Monday's snapshot, then the snapshots from Tuesday and Wednesday must be deleted first.

You can use Remote Copy to roll a volume back to an old snapshot without losing the interim snapshots. Because Remote Copy creates two sets of snapshots—primary and remote snapshots—you can roll a volume back to a snapshot and still retain the other set of snapshots.

# Configuration for nondestructive rollback

To use remote snapshots for nondestructive rollback:

- Create a remote snapshot schedule.
- In the schedule, specify the same retention policy for the primary and remote snapshots. This ensures that you have copies of the same number of snapshots in your primary and remote locations. Any snapshots destroyed during rollback of one volume remain intact on the other volume. See an illustration of a nondestructive rollback configuration in Figure 18 on page 53.

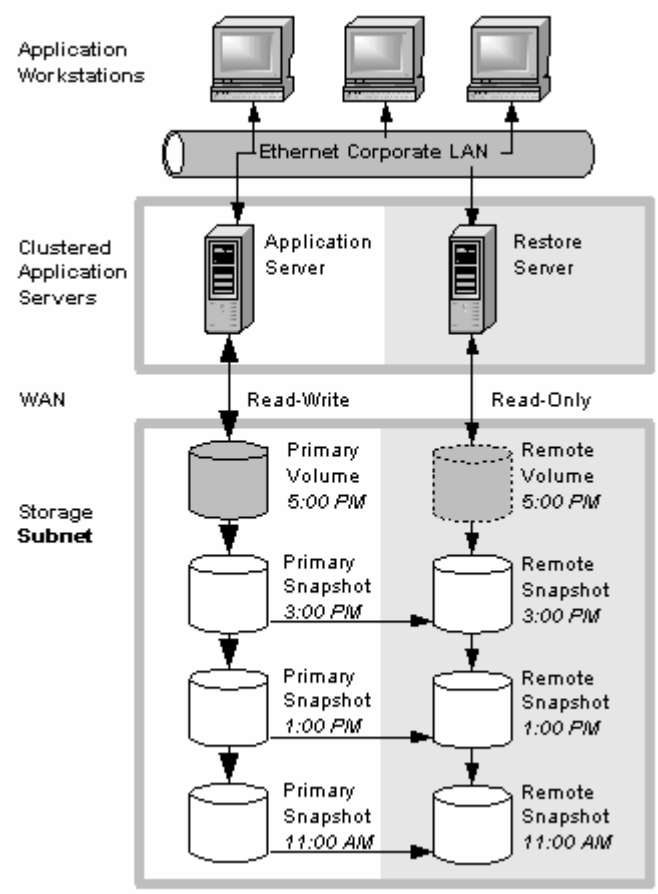


**Figure 18 nondestructive rollback example**

## How this configuration works for nondestructive rollback

You can choose to roll back either the primary snapshot or the remote snapshot. Rolling back one of the snapshots requires that you delete more recent snapshots of that volume. The other volume retains the full set of snapshots. You can continue to make snapshots even though one side was rolled back and the other side was not.

When deciding whether to roll back the primary or remote volume, consider the following:

- When you roll back the primary snapshot to a primary volume, any applications accessing the primary volume will no longer have access to the most current data (because the primary volume has been rolled back to a previous state). If the primary volume must be synchronized with other volumes accessed by the same application, consider rolling back the remote volume instead. Figure 19 on page 54 shows rollback of the primary snapshot while leaving the remote snapshots intact.

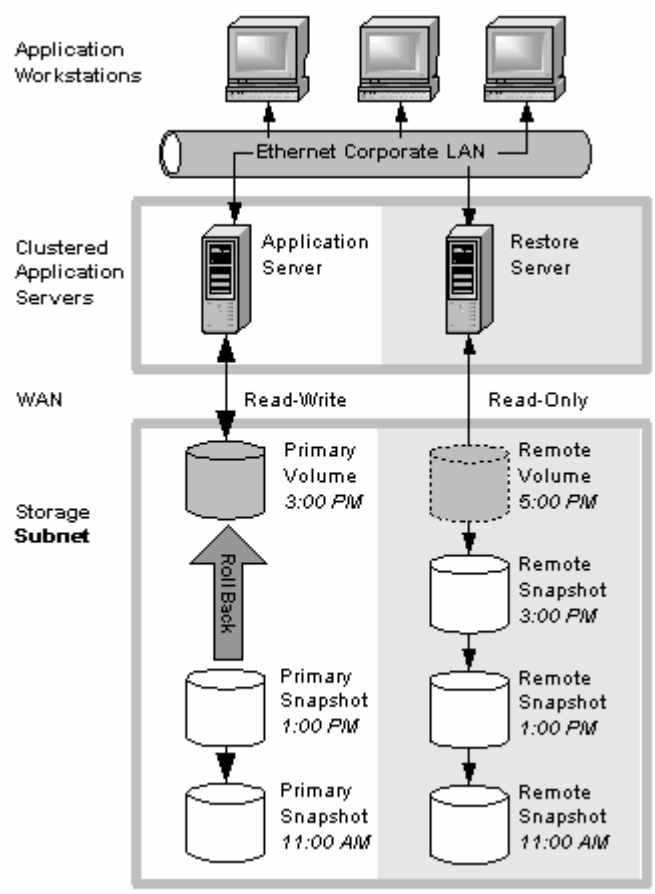


**Figure 19 nondestructive rollback from the primary snapshot**

- To roll back the remote snapshot, you must first make the remote volume into a primary volume. This stops scheduled creation of remote snapshots, which may jeopardize your high availability, disaster recovery, or routine backup strategies. Figure 20 on page 55 shows rollback of the remote snapshot.

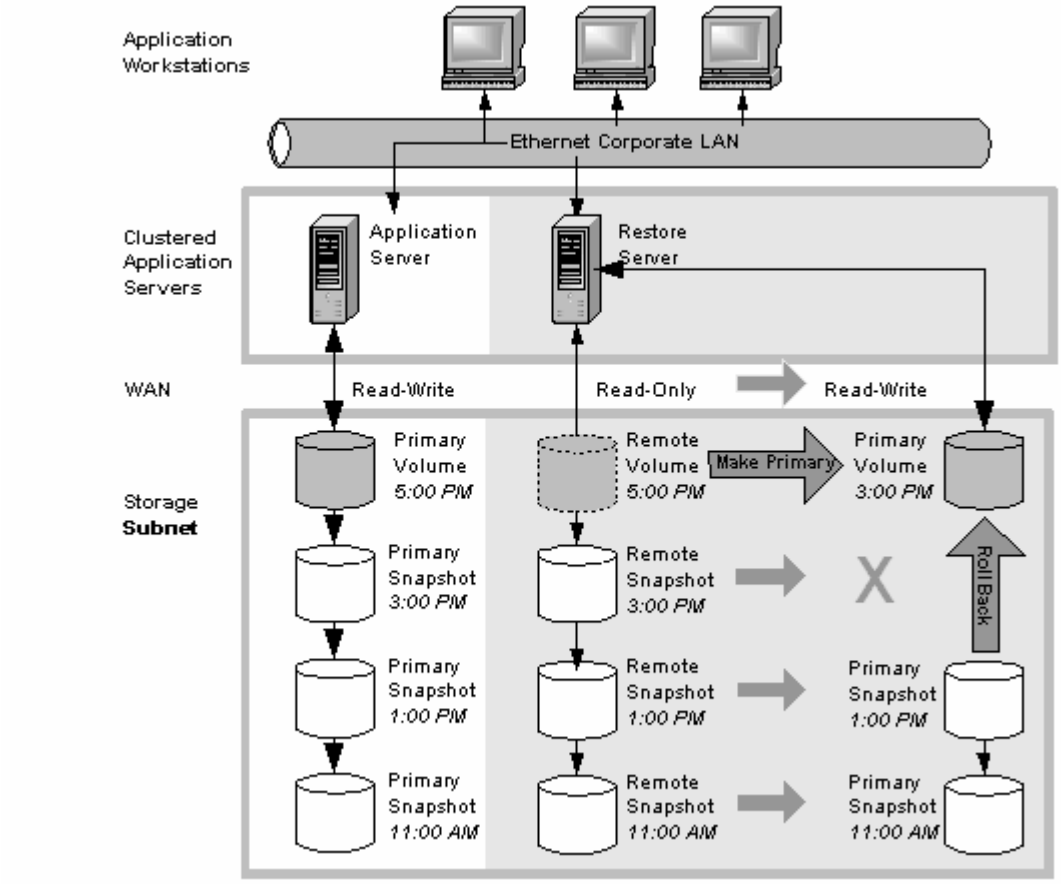


**Figure 20 nondestructive rollback from the remote snapshot**

## Best practices

### Roll back the primary snapshot and keep the remote snapshots as a backup.

To ensure that Remote Copy continues to operate, roll back the primary volume as follows:

1. Preserve the current state of the primary volume that you want to roll back by creating a one-time (manual) remote snapshot of it.
2. Roll back the volume.

   Before roll back, scheduled remote snapshots fail. After the primary volume is rolled back, scheduled creation of remote copies will resume correctly.

   Completed remote snapshots remain intact.

# Using Remote Copy for data migration or cloning

Remote Copy allows migration of data from one application server to another without interrupting the production application server. This capability supports a number of uses such as data mining or content distribution.

## Achieving data migration

You can use Remote Copy to make a complete copy (clone) of one or more volumes without interrupting access to the original volumes. This type of data migration allows you to copy an entire data set for use by a new application or workgroup.

To copy data from one location to another, simply create a one-time remote snapshot of the volume. To make the remote snapshot a read/write volume, make it into a primary volume.

## Configuration for data migration

To make a copy of a volume in a remote location, configure a cluster of storage systems in the remote location with enough space to accommodate the volume. See Figure 21 on page 56 for an example configuration.

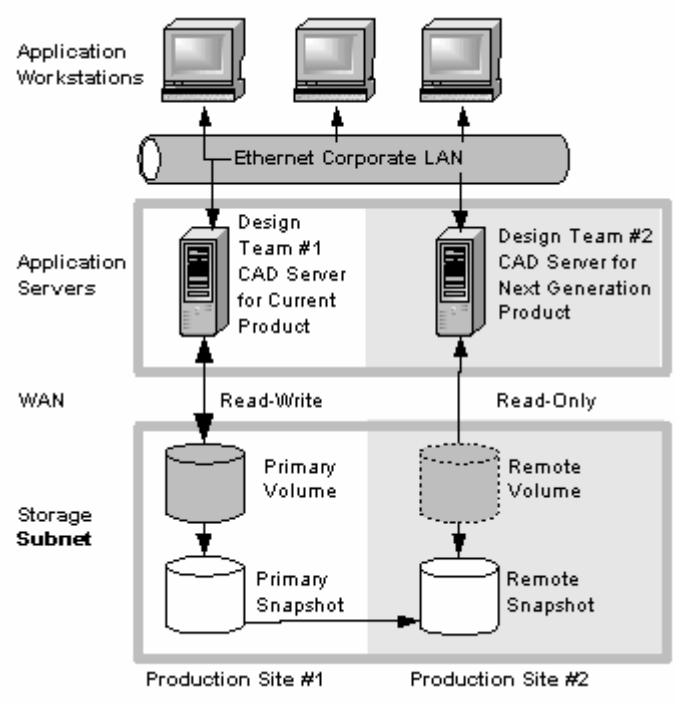


**Figure 21 Data migration example configuration**

## How this configuration works for data migration

Suppose you want to create a complete copy of a volume for an application to use in a different location.

1. Configure a cluster of storage systems in the new location to contain the copied volume.
2. Create a one-time remote snapshot of the volume onto the cluster in the new location.

   If your application server uses multiple volumes that must be in sync, use a script to quiesce the application before creating remote snapshots.

   [Optional] You can create regular one-time snapshots and use remote copy to move the snapshots to the remote cluster at your convenience.
3. On the cluster in the new location, make the remote volume into a primary volume.

4. Configure the application server in the new location to access the new primary volume.

   Figure 22 on page 57 shows the migration of data after converting a remote volume into a primary volume.

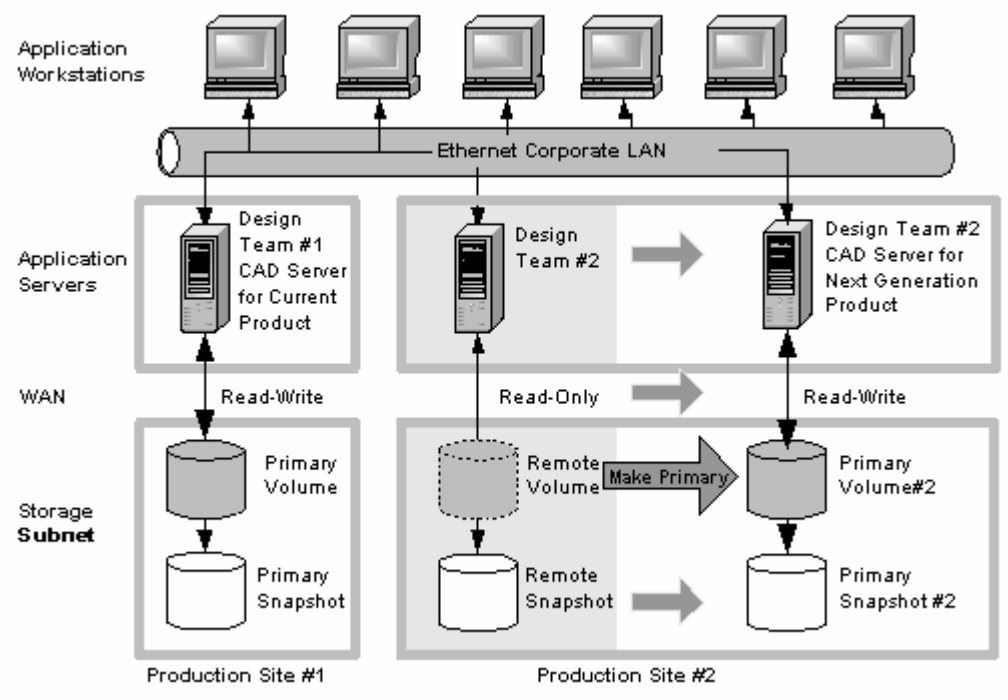


**Figure 22 Configuration after data migration**

# 4 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.

## HP Insight Remote Support Software

HP strongly recommends that you install HP Insight Remote Support software to complete the installation or upgrade of your product and to enable enhanced delivery of your HP Warranty, HP Care Pack Service or HP contractual support agreement. HP Insight Remote Support supplements your monitoring, 24x7 to ensure maximum system availability by providing intelligent event diagnosis, and automatic, secure submission of hardware event notifications to HP, which will initiate a fast and accurate resolution, based on your product's service level. Notifications may be sent to your authorized HP Channel Partner for on-site service, if configured and available in your country. The software is available in two variants:

- HP Insight Remote Support Standard: This software supports server and storage devices and is optimized for environments with 1-50 servers. Ideal for customers who can benefit from proactive notification, but do not need proactive service delivery and integration with a management platform.
- HP Insight Remote Support Advanced: This software provides comprehensive remote monitoring and proactive service support for nearly all HP servers, storage, network, and SAN environments, plus selected non-HP servers that have a support obligation with HP. It is integrated with HP Systems Insight Manager. A dedicated server is recommended to host both HP Systems Insight Manager and HP Insight Remote Support Advanced.

Details for both versions are available at: http://h18004.www1.hp.com/products/servers/management/insight-remote-support/overview.html

To download the software, go to Software Depot:

https://h20392.www2.hp.com/portal/swdepot/index.do.

Select Insight Remote Support from the menu on the right.

# New and changed information in this edition

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

- The following information has been updated:
  - P4000 SAN Solution software and user manuals have been rebranded

# Related information

The following documents [and websites] provide related information:

- *HP StorageWorks P4000 SAN Solution user guide*
- *HP StorageWorks P4000 Multi-Site HA/DR Solution Pack user guide*
- *CLIQ — The SAN/iQ Command-Line Interface 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
- http://www.hp.com/storage/whitepapers

# Typographic conventions

**Table 8 Document conventions**

| Convention | Element |
|---|---|
| Blue text: Table 8 | Cross-reference links and e-mail addresses |
| Blue, underlined text: http://www.hp.com | Website addresses |
| **Bold** text | • Keys that are pressed<br>• Text typed into a GUI element, such as a box<br>• GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes |
| *Italic* text | Text emphasis |

| Convention | Element |
|---|---|
| `Monospace` text | • File and directory names<br>• System output<br>• Code<br>• Commands, their arguments, and argument values |
| *`Monospace, italic`* text | • Code variables<br>• Command variables |
| `Monospace, bold` text | Emphasized monospace text |

⚠ WARNING!

Indicates that failure to follow directions could result in bodily harm or death.

△ CAUTION:

Indicates that failure to follow directions could result in damage to equipment or data.

IMPORTANT:

Provides clarifying information or specific instructions.

NOTE:

Provides additional information.

TIP:

Provides helpful hints and shortcuts.

# Customer self repair

HP customer self repair (CSR) programs allow you to repair your StorageWorks product. If a CSR part needs replacing, HP ships the part directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your HP-authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider, or see the CSR website:

http://www.hp.com/go/selfrepair

This product has no customer replaceable components.

# 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 StorageWorks P4000 SAN Solution.

| | |
|---|---|
| **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 systems on the subnet the CMC is connected to. Any storage systems 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 system 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 related to 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 systems that create the storage pool from which you create volumes. |
| **CMC** | Centralized Management Console. See HP StorageWorks P4000 Centralized Management Console. |

| | |
|---|---|
| **communication mode** | The unicast communication among storage systems 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 systems, and iSCSI sessions in the HP StorageWorks P4000 SAN 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:<br>• Active - on and participating in RAID<br>• Uninitialized or Inactive - On but not participating in RAID<br>• Off or Missing - Not on<br>• DMA Off - disk unavailable due to faulty hardware or improperly seated in the chassis |
| **DSM** | Device Specific Module |
| **DSM for MPIO** | The HP StorageWorks P4000 DSM for MPIO 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 system 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 system** | When using Repair Storage System, a "ghost" storage system acts as a placeholder in the cluster, keeping the cluster intact, while you repair or replace the storage system. |

**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 system, its drives, and configuration.

**hostname** The hostname on a storage system is the user-definable name that displays below the storage system icon in the network window. It is also visible when the users browse the network.

**HP StorageWorks P4000 Centralized Management Console** Management interface for the SAN/iQ software.

**ID LED** LED lights on the physical storage system so that you can find that system in a rack.

**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 systems in a cluster.

**license keys** A license key registers a storage system for add-on applications. Each storage system 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 system are stored both locally on the storage system 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 systems which serves as the container within which you cluster storage systems and create volumes for storage.

**managers** Manager software runs on storage systems within a management group. You start managers on designated storage systems to govern the activity of all of the storage systems in the group.

**MIB** The Management Information Base provides SNMP read-only access to the storage system.

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

- Contain storage systems that reside in two or more sites
- Contain storage systems 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 systems in the cluster. Network RAID-10, 10+1 or 10+2 is required to protect data in an HP P4000 SAN solution. |
| **network window** | Graphically depicts the status of each storage system. Storage systems on the network are either available or part of a management group. |
| **NTP** | Network Time Protocol |
| **parity** | In RAID 5, redundant information is stored as parity distributed across the disks. Parity allows the storage system 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 StorageWorks P4000 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 related to 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:<br>• RAID 0 - data striped across disk set<br>• RAID 1 - data mirrored from one disk onto a second disk<br>• RAID 10 - mirrored sets of RAID 1 disks |

- RAID 5 - data blocks are distributed across all disks in a RAID set. Redundant information is stored as parity distributed across the disks.
- RAID 50 - mirrored sets of RAID 5 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 system:

- 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 system. The storage system is offline and flashes red in the network window.

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

**remote copy pair** The primary volume and its related 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 system** Creates a placeholder in the cluster, in the form of a “ghost” storage system, that keeps the cluster intact while you remove the storage system to replace a disk or replace the storage system 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 system, or remove a storage system. 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 system goes down, and writes continue to a second storage system, and the original store comes back up, the original storage system needs to recoup the exact data captured by the second storage system. |
| **rolling back** | Replaces the original volume with a read/write copy of a selected snapshot. Starting with release 8.0, he new volume retains the same name. |
| **SAN/iQ interface** | When you initially set up a storage system 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 systems are installed. Multi-Site SAN configurations have multiple sites with storage systems 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 StorageWorks 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:<br>• Cannot contain storage systems that are designated to reside in a site.<br>• Cannot contain storage systems that span subnets.<br>• 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 system in a cluster becomes unavailable, servers can still access the volume through the other storage systems 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 system until it is needed to regain quorum. |
| **volume** | A logical entity that is made up of storage on one or more storage systems. 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 StorageWorks P4000 VSS Provider is the hardware provider that supports the Volume Shadow Copy Service on the HP P4000 SAN Solution.

**writable space** *See* temporary space

# Index

## S



## T



## U



## V

## W