# LeftHand Networks® Remote Copy User Guide

LeftHand Networks, Inc.
2580 55th Street
Boulder CO 80301
(303) 449-4100
www.LeftHandNetworks.com

This manual is designed to provide accurate and authoritative information in regard to the use of the SAN/iQ Centralized Management Console, and the SAN/iQ software, proprietary products of LeftHand Networks, Inc. Information in this document is subject to change without notice and does not represent a commitment on the part of LeftHand Networks, Inc.

Part # 26200001813 Rev. A

# Table Of Contents

# Understanding and Planning Remote Copy

1

## Remote Copy Overview

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.

## Topics Covered in this Chapter

- "How Remote Copy Works"
- "Planning for Remote Copy"
- "Using Schedules for Remote Copy"

## How Remote Copy Works

Remote Copy uses the existing volume and snapshot features along 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.

## Glossary for Remote Copy

The following terminology is used in describing the components and processes involved in Remote Copy.

Table 1.1. Remote Copy Glossary

| Term | Definition |
|---|---|
| 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. |
| 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. |
| 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:<br>• in the same cluster in the same management group<br>• in a different cluster in a different management group<br>• in a different cluster in the same management group. |
| Remote Snapshot | An identical copy of a primary snapshot. The remote snapshot is located on the same cluster as the remote volume. |
| Remote Copy Pair | The primary volume and its associated 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. |
| Acting Primary Volume | The remote volume, when it assumes the role of the primary volume in a failover scenario. |
| Original Primary Volume | The primary volume that fails and then is returned to service. |
| Failback | After failover, the process by which you restore the primary volume and turn the acting primary back into a remote volume. |
| 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. |

Table 1.1. Remote Copy Glossary (Continued)

| Term | Definition |
|---|---|
| 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.<br>Also, you can manually synchronize if necessary to include data that is on the remote volume but not the primary. |
| 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. |

# How Remote Copy Works

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

1. At the production location, you create a snapshot of the primary volume — this is called the primary snapshot.
2. You create a remote volume at the remote location. Then you create a remote copy of the primary snapshot to the remote volume.
3. The system copies data from the primary snapshot to the remote snapshot.

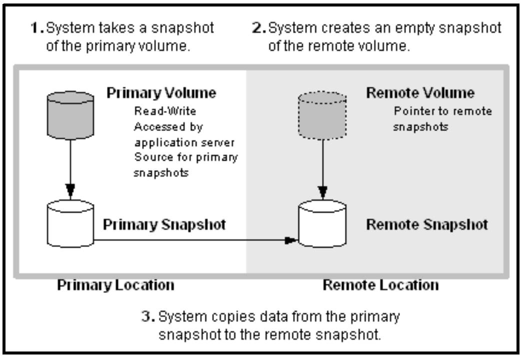
1. System takes a snapshot of the primary volume.
2. System creates an empty snapshot of the remote volume.
Primary Volume
Read-Write
Accessed by application server
Source for primary snapshots
Remote Volume
Pointer to remote snapshots
Primary Snapshot
Remote Snapshot
Primary Location
Remote Location
3. System copies data from the primary snapshot to the remote snapshot.

Figure 1.1. Basic flow of Remote Copy

**Note:** Both primary and completed remote snapshots are the same as regular snapshots. See the chapter “Working with Snapshots” in the *LeftHand SAN User Manual*.

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

## Graphical Representations of Remote Copy

The 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, as illustrated in Figure 1.2. The pages move in the direction of the data flow from primary to remote snapshot.

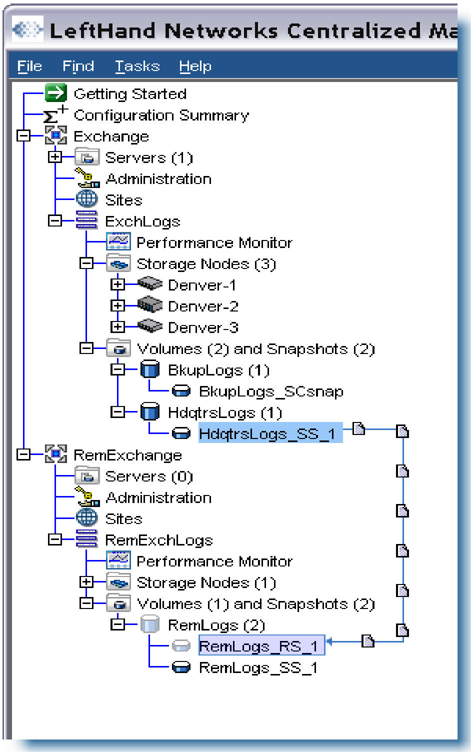

Figure 1.2. Icons depicting the primary snapshot copying to the remote snapshot

## Graphical Legend for Remote Copy Icons

The graphical legend available from the Help menu depicts the icons associated with Remote Copy. Figure 1.3 displays the Remote Copy states icons from the graphical legend.

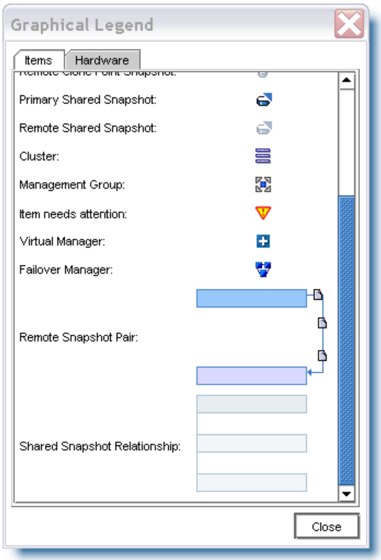

Figure 1.3. Icons for Remote Copy in the Graphical Legends window

## Remote Copy and Volume Replication

Remote Copy is asynchronous replication of data. Volume replication is synchronous replication. Volume replication is described in detail in the *LeftHand SAN User Manual* in the chapter, “Working with Volumes.” Using synchronous volume replication on multiple storage nodes within a cluster in combination with asynchronous Remote Copy on a different cluster of storage nodes creates a robust, high-availability configuration.

## Uses for Remote Copy

Review Table 1.2 to see common applications for the Remote Copy application.

Table 1.2. Uses for Remote Copy

| Use Remote Copy for | How It Works |
|---|---|
| Business continuance/ disaster recovery | Using Remote Copy, store remote snapshots on a machine geographically separate. The remote snapshots remain available in the event of a site or system failure at the primary site. |
| Off-site backup and recovery | Using Remote Copy, eliminate 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 | Using Remote Copy, make a complete copy of one or more volumes without interrupting access to the original volumes. Move the copy of the volume to the location where it is needed. |
| Volume clone | Using Remote Copy, create 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. Hence, 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, snapshot, and storage node level. Review Table 1.3 for common configurations at various levels.

Table 1.3. Remote Copy, SAN/iQ, and Storage Nodes

| 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. |
| 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 are copied from the primary snapshot to the remote snapshot, the remote snapshot behaves as a regular snapshot. |
| Storage Nodes | • Active monitoring of each storage node notifies you when copies complete or fail. Active monitoring 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. |

## Planning the Remote Snapshot

To create a remote snapshot, meet these prerequisites:

- 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.
- Have 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 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 19.

# Using Schedules for Remote Copy

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

## Planning the Remote Copy Schedule

Planning is critical. These 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 4 hours, you should ensure that the time to copy that snapshot to the remote location is less than 4 hours.

Test the time required for copying a snapshot. One way to check the time required to copy a snapshot is to run a test of the actual process. In the test, you take 2 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. Because you create the second remote snapshot after the time interval you intend to schedule, the copy time for the second remote snapshot is more representative of the actual time required for copying subsequent remote snapshots.

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

   Be sure to check the remote bandwidth setting for the other for the management group with 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? See the chapter “Provisioning Storage” in the *LeftHand SAN User Manual* for information about managing capacity.

If the cluster does not have sufficient space available, the remote snapshot appears in the CMC and it flashes red. On the Details tab of the remote snapshot, the status says “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 2 primary snapshots and 5 remote snapshots. The number of snapshots retained refers to completed snapshots. Take the following characteristics 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 more frequently than the remote copy process can keep up with. 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 2 primary and 5 remote snapshots, the software may retain up to 4 primary and 7 remote snapshots for a period of time. Table 1.4 shows the maximum retained snapshots with respect to a specific retention policy.

Table 1.4. Snapshot retention policy and maximum number of retained snapshots

| **Scheduled Remote Snapshot Retention Policy** | **Maximum Number of Snapshots Retained** |
|---|---|
| *n* of primary snapshots<br>*x* of remote snapshots | *n* + 2 of primary snapshots<br>*x* + 2 of remote snapshots |
| *n* of hours for primary snapshots<br>*x* of hours for remote snapshots | *n* + 2 primary snapshots older than *n*<br>*x* + 2 remote snapshots older than *x* |
| *n* of days for primary snapshots<br>*x* of days for remote snapshots | *n* + 2 primary snapshots older than *n*<br>*x* + 2 remote snapshots older than *x* |
| *n* of weeks for primary snapshots<br>*x* of weeks for remote snapshots | *n* + 2 primary snapshots older than *n*<br>*x* + 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 2 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 1.5 to help plan scheduled remote snapshots.

Table 1.5. Scheduled Remote Copy Planning Checklist

| Configuration Category | Characteristic |
|---|---|
| **Scheduled Snapshot** | |
| Start Time | • Start date (mm/dd/yyyy)<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. |

Table 1.5. Scheduled Remote Copy Planning Checklist (Continued)

| Configuration Category | Characteristic |
|---|---|
| **Primary Setup** | |
| Retention | Retain either<br>• Maximum number of snapshots (#)<br>• Set period of time (minutes, hours, days or weeks) |
| **Remote Setup** | |
| Management Group | The management group to contain the remote snapshot |
| Volume | The remote volume for the remote snapshots |
| Retention | Retain either<br>• Maximum number of snapshots (#). This number equals completed snapshots only. In-progress snapshots take additional space on the cluster while they are being copied. Also, the system will not delete the last fully synchronized snapshot. For space calculations, figure N+2 with N=maximum number of snapshots.<br>• Set period of time (minutes, hours, days or weeks) |

# Using Remote Copy

# 2

## Remote Copy Overview

This chapter provides instructions for registering, configuring, and using Remote Copy for business continuance, backup and recovery, and failover.

For information about how Remote Copy works and how to plan capacity for Remote Copy, see Chapter 1, “Understanding and Planning Remote Copy”.

## 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.

### 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.

Creating a remote snapshot involves these main steps:

- Log in to the primary and the remote management group.
- Create a primary snapshot of the primary volume manually. When doing a schedule to remote snapshot 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 management groups, 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, you only log in to the one 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 select New Remote Snapshot.

   The New Remote Snapshot window opens, as shown in Figure 2.1.

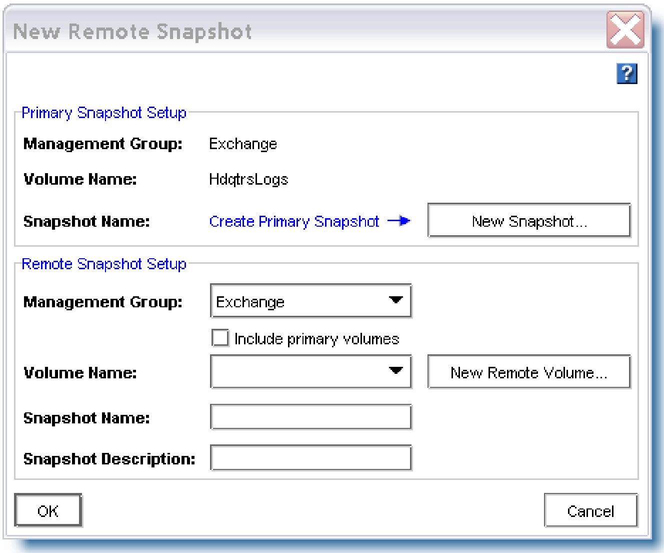

**Management Group:** Exchange

**Volume Name:** HdqtrsLogs

**Snapshot Name:** Create Primary Snapshot →

**Management Group:** Exchange

□ Include primary volumes

**Volume Name:**

**Snapshot Name:**

**Snapshot Description:**

Figure 2.1. Creating a new remote snapshot

## Creating the Primary Snapshot

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

   The New Snapshot window opens, shown in Figure 2.2.

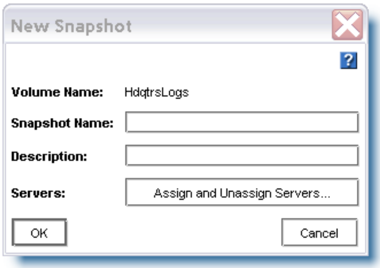

**Volume Name:** HdqtrsLogs

**Snapshot Name:**

**Description:**

**Servers:**

Figure 2.2. Creating a new primary snapshot

2. Enter a name for the snapshot.

**Tip:** Make the beginning of volume and snapshot names meaningful, for example, “Snap1Exchg_03.”

3. (Optional) Type in a description of the snapshot.
4. Click OK to return to the New Remote Snapshot window.

   The information for the primary snapshot is filled in, as shown in Figure 2.3. That is, the text for the field Snapshot Name has changed:

   - From ‘Create Primary Snapshot’
   - To ‘HdqtrsLogs_SS_1’

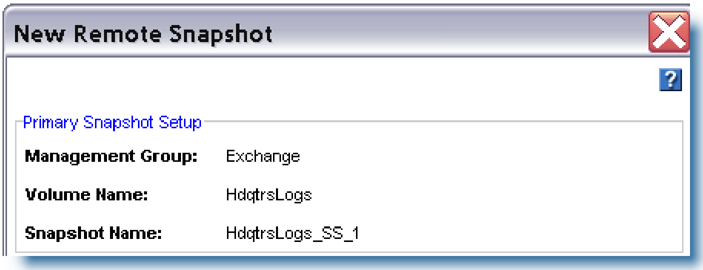

New Remote Snapshot

Primary Snapshot Setup

Management Group: Exchange

Volume Name: HdqtrsLogs

Snapshot Name: HdqtrsLogs_SS_1

Figure 2.3. New primary snapshot created

In the Remote Snapshot Setup box, use the drop-down lists and select the remote site management group and volume.

The Management Groups, Clusters and Volumes wizard is available at this point if you need to create the remote management group.

5. In the Snapshot Name field, type in the name for this remote snapshot.
6. (Optional) Type in a description for the remote snapshot.
7. Click OK in the New Remote Snapshot window.

   The remote copy of the primary snapshot to the remote volume begins, as shown in Figure 2.4.

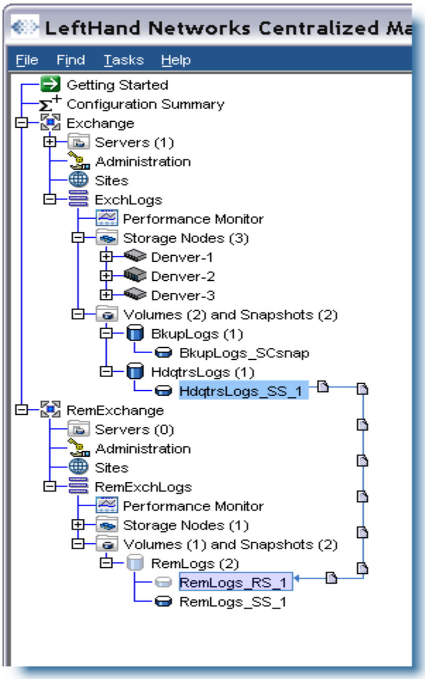

Figure 2.4. Remote copy in progress

## Creating a Remote Volume

Create a remote volume by any of the following methods:

- Designate an existing primary volume as a remote volume. See “Designating an existing volume as a remote volume” on page 20.
- Create a new remote volume manually. See “Creating a new remote volume manually” on page 20.
- Create a new remote volume during creation of a remote snapshot. See “Creating a remote volume on the fly” on page 20.
- Use the Management Groups, Clusters, and Volumes wizard in Getting Started. See Chapter 1, “Getting Started.” in the *LeftHand SAN User Manual* for details on working through the wizards.

### Designating an existing volume as a remote volume

Selecting an existing volume to become a remote volume causes the following:

- A scheduled snapshot is created of the volume, and then
- The volume becomes a 0-byte remote volume.

See “Making a Primary Volume Into a Remote Volume” on page 43.

### Creating a new remote volume manually

Create a remote volume as you would any other volume. Be sure to choose the storage nodes 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. You create a remote volume at the remote site to receive the snapshot. Then, either take the snapshot and remote copy it, or create the schedule to take remote snapshots.

See the section on “Creating a Volume” in the *LeftHand SAN User Manual***,** the chapter “Working With Volumes.”

### Creating a remote volume on the fly

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.

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

2. Click New Volume.

   The Management Groups, Clusters, and Volumes wizard opens.

   For specific help, see Chapter 1, “Getting Started.” in the *LeftHand SAN User Manual* for details on working through the wizards.

   The information you specify in the wizard fills in the New Remote Snapshot window when you exit the wizard.

3. (Optional) Type in a description of the remote snapshot and click OK.

   The remote copy may take some time.

## What the System Does

The system creates the remote snapshot in the cluster that contains the remote volume.

The system 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, as shown in Figure 2.5.

**Note:** If you create a remote snapshot of a volume with a remote snapshot still in progress, the second remote snapshot does not begin copying until the first remote snapshot is complete.

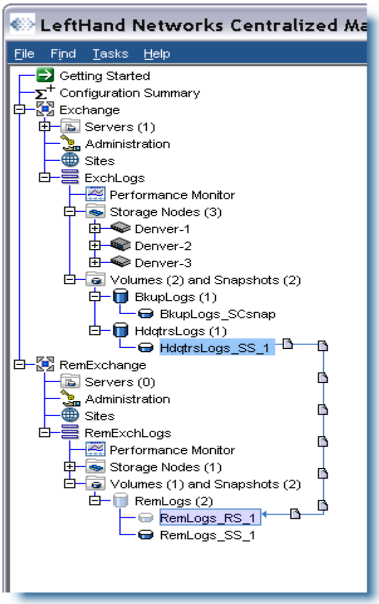

Figure 2.5. Viewing the remote snapshot

## 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 have accumulated much data in the primary site and your hardware is already installed in the remote site.

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

**Use the storage nodes intended for the remote site to configure the remote management group on-site and copy data locally. Then, ship the remote storage nodes to the remote site.**

Use this method if you initially have all the storage nodes 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 nodes for the remote site and install the remote management group just as you configured it in the primary site.

   Allow adequate time between the arrival of the storage nodes and the first remote copy operation.

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

**Use the PrimeSync feature of Remote Copy to configure a temporary management group, create the first copy locally, then ship the temporary storage node and again 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 nodes to create a new temporary (PrimeSync) management group, cluster and volume.
2. **Make a remote snapshot of** the primary Site A volume to the temporary PrimeSync management group over Gigabit Ethernet. See “Creating a Remote Snapshot” on page 15.
3. Ship the storage nodes to the remote Site B. Power them on and discover them in the CMC to display the temporary PrimeSync management group, cluster and volume.
4. Copy the remote snapshot from the temporary PrimeSync management group to the existing remote Site B management group.

5. Disassociate the temporary PrimeSync management group from the remote Site B management group. For detailed instructions, see “Disassociating Remote Management Groups” on page 60.
6. Delete the temporary PrimeSync management group.
7. Now set up the desired Remote Copy relationship, such as configuring a schedule to remote snapshot the volume from the primary Site A to remote Site B management group.

Note

**Note:** Be sure you use the first initial snapshot, used for the temporary PrimeSync management group copy, to create the second remote copy or the schedule to remote snapshot the volume. You are now setting up the remote copy that goes from primary site A directly to remote site B and keeps that relationship going forward.

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.

For more information on PrimeSync, look for the “Application Note: SAN/iQ Remote Copy PrimeSync—Creating Initial Copy” in the Customer Resource Center.

## 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 the cluster for which you want to view the list of remote snapshots.
2. Click the Remote Snapshot tab to bring it to the front, shown in Figure 2.6.

   The report in the tab window lists management groups and all the snapshots. The other columns show status information about the remote snapshots as described in detail in “Monitoring Remote Snapshots” on page 30.

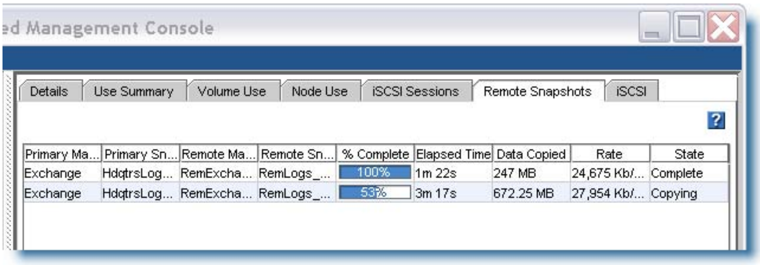

| Primary Ma... | Primary Sn... | Remote Ma... | Remote Sn... | % Complete | Elapsed Time | Data Copied | Rate | State |
|---|---|---|---|---|---|---|---|---|
| Exchange | HdqtrsLog... | RemExcha... | RemLogs_... | 100% | 1m 22s | 247 MB | 24,675 Kb/... | Complete |
| Exchange | HdqtrsLog... | RemExcha... | RemLogs_... | 53% | 3m 17s | 672.25 MB | 27,954 Kb/... | Copying |

Figure 2.6. Viewing the list of remote snapshots

## Setting the Remote Bandwidth

The remote bandwidth sets the maximum rate for data transfer between management groups. That is, 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—the remote management group. When setting the remote bandwidth, you can choose from a list of common network types, or you can calculate a custom rate, based on your particular requirements.

### Select 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, the speed is the maximum rate at which data will be copied.

**Defaults Setting**

When setting remote bandwidth, selecting Defaults allows you to choose from a list of common network types, as shown in Figure 2.7.

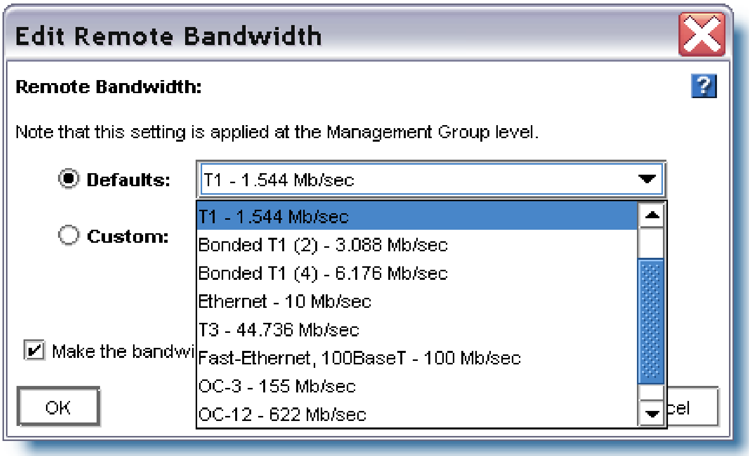

Edit Remote Bandwidth

Remote Bandwidth:

Note that this setting is applied at the Management Group level.

● Defaults: T1 - 1.544 Mb/sec

○ Custom:

☑ Make the bandwi

Figure 2.7. Defaults setting for remote bandwidth

### Custom Setting

The custom setting for remote bandwidth defaults to 32768 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, as shown in Figure 2.8.

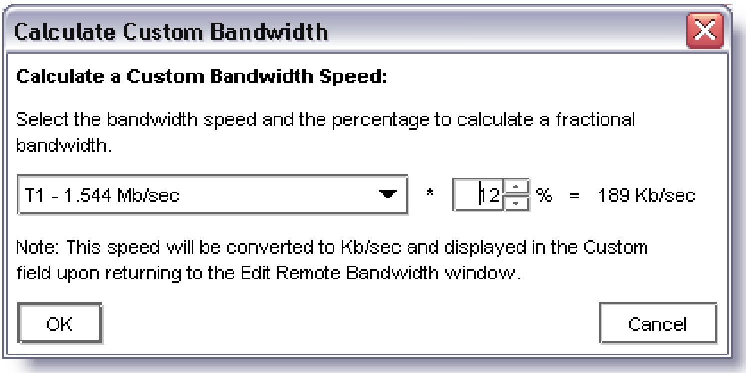

Calculate Custom Bandwidth

Calculate a Custom Bandwidth Speed:

Select the bandwidth speed and the percentage to calculate a fractional bandwidth.

T1 - 1.544 Mb/sec * 12 % = 189 Kb/sec

Note: This speed will be converted to Kb/sec and displayed in the Custom field upon returning to the Edit Remote Bandwidth window.

Figure 2.8. 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.

## To Set the Bandwidth

1. In the navigation window, select the management group, either the remote or the primary one.
2. Click Management Group Tasks and select Edit Management Group.

   The Edit Management Group window opens, as shown in Figure 2.9.
3. Select the remote management group, remote or primary.

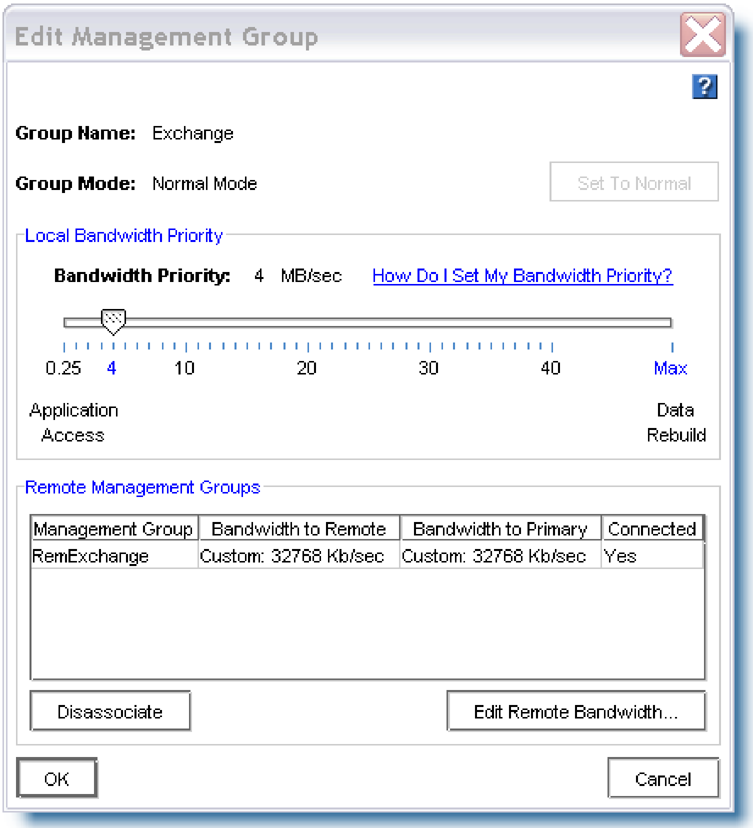

**Group Name:** Exchange

**Group Mode:** Normal Mode

**Bandwidth Priority:** 4 MB/sec

| Management Group | Bandwidth to Remote | Bandwidth to Primary | Connected |
|---|---|---|---|
| RemExchange | Custom: 32768 Kb/sec | Custom: 32768 Kb/sec | Yes |

Figure 2.9. Editing a remote management group

4. Click Edit Remote Bandwidth.

   The Edit Remote Bandwidth window opens, as shown in Figure 2.10.

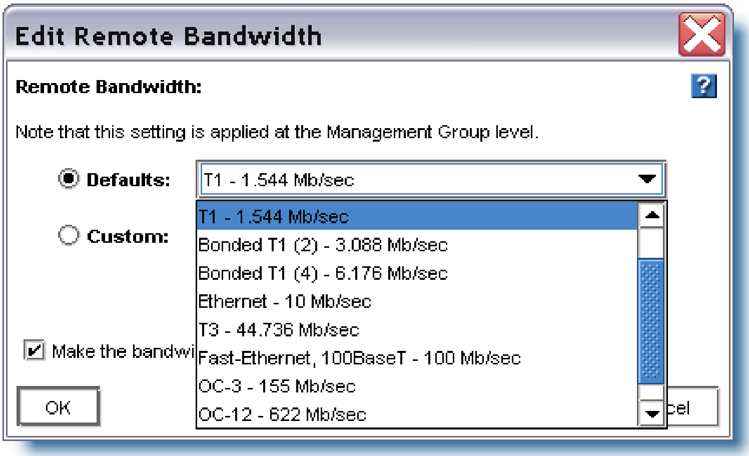

Edit Remote Bandwidth

Remote Bandwidth:

Note that this setting is applied at the Management Group level.

● Defaults: T1 - 1.544 Mb/sec

○ Custom:

☑ Make the bandwi

Figure 2.10. Editing the remote bandwidth

5. Change the bandwidth setting as desired.

## Canceling a Remote Snapshot

When you cancel a remote snapshot as it 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 select Cancel Remote Snapshot.

   A confirmation message opens.
5. Click OK.

## Editing a Remote Snapshot

You can 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 select Edit Snapshot.

   The Edit Snapshot window opens, shown in Figure 2.11.

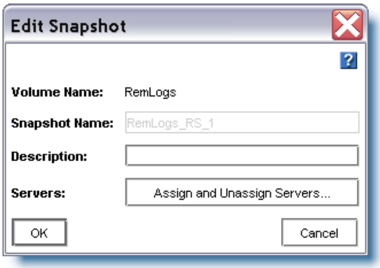

Edit Snapshot

Volume Name: RemLogs

Snapshot Name: RemLogs_RS_1

Description:

Servers:

Figure 2.11. Editing a remote snapshot

4. Change the desired information and click OK.

## Deleting 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 Snapshot Tasks and select Delete Snapshot from the menu.

   A confirmation message opens.
4. Click OK and then click Delete Snapshot in the next confirmation window.

# Monitoring Remote Snapshots

Information for monitoring remote snapshots is available from multiple sources. Active monitoring provides you with the capability to configure alerts that you view in the alert window as well as receiving alerts as emails and through SNMP traps. The Alert window also provides real time monitoring information for remote snapshots when you are logged in to the CMC.

## Monitoring Remote Snapshots Details from the Tab Window

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

### Viewing Information in the Remote Snapshot Tab

The Remote Snapshots tab displays a list of remote snapshots connected with a selected item in the navigation window. For example, if you select a management group, the Remote Snapshots tab displays the list of remote snapshots associated with that management group. 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 to bring it to the front, shown in Figure 2.12.

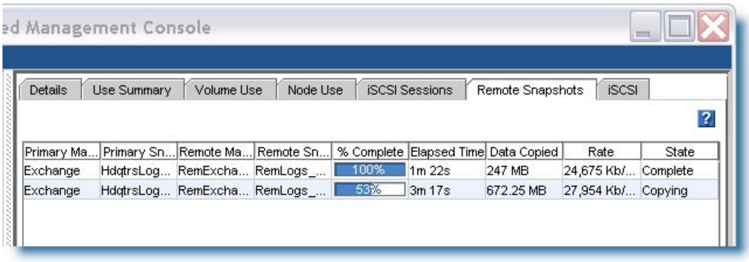

| Primary Ma... | Primary Sn... | Remote Ma... | Remote Sn... | % Complete | Elapsed Time | Data Copied | Rate | State |
|---|---|---|---|---|---|---|---|---|
| Exchange | HdqtrsLog... | RemExcha... | RemLogs_... | 100% | 1m 22s | 247 MB | 24,675 Kb/... | Complete |
| Exchange | HdqtrsLog... | RemExcha... | RemLogs_... | 53% | 3m 17s | 672.25 MB | 27,954 Kb/... | Copying |

Figure 2.12. Viewing remote snapshot details in the Remote Snapshots tab

You may want to check the remote snapshot details for this information:

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

## Viewing 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 to bring it to the front.
2. Select a remote snapshot from the list.
3. Click Remote Snapshot Tasks and select View Remote Snapshot Details.

   The Remote Snapshot Details window opens, as shown in Figure 2.13.

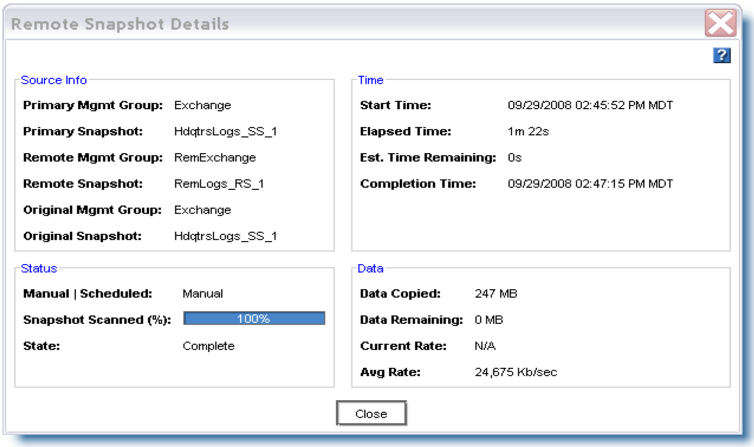

Source Info

Primary Mgmt Group: Exchange
Primary Snapshot: HdqtrsLogs_SS_1
Remote Mgmt Group: RemExchange
Remote Snapshot: RemLogs_RS_1
Original Mgmt Group: Exchange
Original Snapshot: HdqtrsLogs_SS_1

Time

Start Time: 09/29/2008 02:45:52 PM MDT
Elapsed Time: 1m 22s
Est. Time Remaining: 0s
Completion Time: 09/29/2008 02:47:15 PM MDT

Status

Manual | Scheduled: Manual
Snapshot Scanned (%): 100%
State: Complete

Data

Data Copied: 247 MB
Data Remaining: 0 MB
Current Rate: N/A
Avg Rate: 24,675 Kb/sec

Figure 2.13. Viewing remote snapshot details

During the remote copy process, the Remote Snapshot Details window reports current data for the statistics. When the copy is completed, the statistics show summary data. Figure 2.13 shows a completed remote copy. Table 2.1 lists the values for the statistics reported in the Remote Snapshot Details window.

Table 2.1. Values for Remote Snapshot Details window

| Statistic | Values |
|---|---|
| **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. Used with PrimeSync feature. |
| Original Snapshot | The first version of the snapshot from which the first copy was created. Used with PrimeSync feature. |

Table 2.1. Values for Remote Snapshot Details window (Continued)

| Statistic | Values |
|---|---|
| **Status** | |
| Manual \| Scheduled | Whether the snapshot was created manually or with a scheduled snapshot. |
| State | Started, Copying, Stalled, Complete<br>Current state of the copy process. |
| Snapshot Scanned (%) | 0-100%<br>Percent of the copy process that is completed. |
| **Time** | |
| Start Time | MM/DD/YYYY HH:MM:SS [AM/PM] Time Zone<br>Date and time copy started |
| Elapsed Time | Xd Xh Xm Xs<br>X = a number and the days, hours, minutes, and seconds the copy has been processing.<br>N/A if not yet available. |
| Est. Time Remaining | Xd Xh Xm Xs<br>X = a number and the days, hours, minutes, and seconds estimated to remain in the copy process.<br>N/A for completed copies or in-progress copies not yet calculated. |
| Completion Time | MM/DD/YYYY HH:MM:SS [AM/PM] Time Zone<br>Date and time copy completed.<br>N/A for in-progress copies. |
| **Data** | |
| Data Copied | MB, GB, or TB<br>Amount of data copied so far in smallest unit size. |
| Data Remaining | MB, GB, or TB<br>Amount of data remaining to be copied in smallest unit size |
| Current Rate | Kb/sec.<br>Current rate of data being copied in Kb/second. This rate is recalculated regularly throughout the remote copy process.<br>N/A if not yet available or completed. |
| Avg. Rate | Kb/sec.<br>Average rate of copy progress. |

You can leave the Remote Snapshot Details window open and monitor the progress of the remote copy. An example of a Remote Snapshot Details window with a remote copy in progress is shown in Figure 2.14.

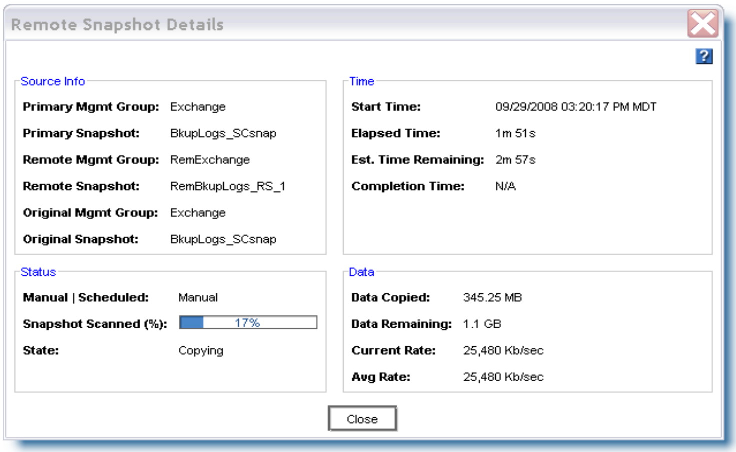

Remote Snapshot Details

Source Info

Primary Mgmt Group: Exchange

Primary Snapshot: BkupLogs_SCsnap

Remote Mgmt Group: RemExchange

Remote Snapshot: RemBkupLogs_RS_1

Original Mgmt Group: Exchange

Original Snapshot: BkupLogs_SCsnap

Time

Start Time: 09/29/2008 03:20:17 PM MDT

Elapsed Time: 1m 51s

Est. Time Remaining: 2m 57s

Completion Time: N/A

Status

Manual | Scheduled: Manual

Snapshot Scanned (%): 17%

State: Copying

Data

Data Copied: 345.25 MB

Data Remaining: 1.1 GB

Current Rate: 25,480 Kb/sec

Avg Rate: 25,480 Kb/sec

Figure 2.14. Viewing remote snapshot details for remote copy in progress

## Configuring Active Monitoring Alerts for Remote Copy

There are four variables for remote snapshots for which you can configure alerts. Notification for these variables automatically displays as default alert messages in the alert window. You can also configure Active Monitoring to receive email notification or for SNMP traps. The Remote Copy variables that are monitored include these:

- Remote Copy status—an alert is generated if the copy fails.
- Remote Copy complete—an alert is generated when the remote copy is complete.
- Remote Copy failovers—an alert is generated when a remote volume is made primary.
- Remote management group status—an alert is generated if the connection to a remote management group changes (disconnects and/or reconnects).

To read about configuring Active Monitoring, see the Reporting chapter of the *LeftHand SAN User Manual*.

# 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 for business continuance and disaster recovery, as well as provide 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 8.

## Best Practices for Scheduling Remote Snapshots

- Create a new remote volume to use with the scheduled remote snapshots.
- If performing daily remote copies, schedule the remote snapshots during off-peak hours. If setting scheduled remote snapshots for multiple volumes, stagger the schedules with at least an hour between start times for best results.
- Use NTP to set all storage nodes in the management group to the same time zone.
- Reset the management group time before creating a new schedule to remote snapshot a volume. For detailed information, see “Resetting the Management Group Time” in the

chapter “Working with Management Groups” in the *LeftHand SAN User Manual*.

## Creating the Schedule

Create the schedule for continuing remote snapshots.

1. In the navigation window, select the primary volume.
2. Click the Schedules tab to bring it to the front.
3. Click Schedule Tasks and select New Schedule to Remote Snapshot a Volume.

   The New Schedule to Remote Snapshot a Volume window opens, shown in Figure 2.15.

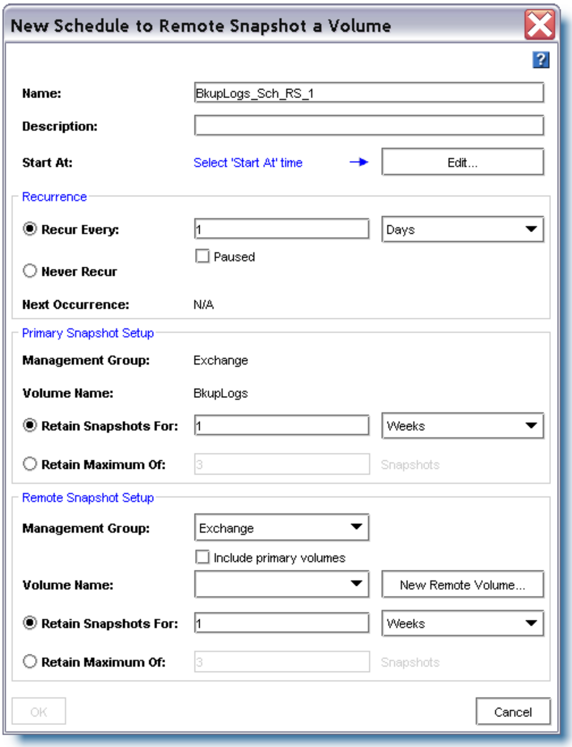

New Schedule to Remote Snapshot a Volume

Name: BkupLogs_Sch_RS_1

Description:

Start At: Select 'Start At' time →

Recurrence

● Recur Every: 1 Days

□ Paused

○ Never Recur

Next Occurrence: N/A

Primary Snapshot Setup

Management Group: Exchange

Volume Name: BkupLogs

● Retain Snapshots For: 1 Weeks

○ Retain Maximum Of: 3 Snapshots

Remote Snapshot Setup

Management Group: Exchange

□ Include primary volumes

Volume Name:

● Retain Snapshots For: 1 Weeks

○ Retain Maximum Of: 3 Snapshots

Figure 2.15. Creating a new schedule for recurring remote snapshots

4. Click Edit and select a ‘Start At’ time.
5. Select a recurrence interval.

   You can schedule a snapshot 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, and up to 200 for all volumes.
7. Select the management group and volume that will hold the *remote* snapshots.

   Log in if you need to.

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

   You can retain up to 50 snapshots for a volume, and up to 200 for all volumes.
9. 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.

## Timing for a Scheduled Remote Snapshot

When you set up a schedule for recurring remote snapshots with the previous procedure, you rely on the time. The time zone displayed in the Schedule to Remote Snapshot a Volume windows is the time zone of the storage node through which you first logged in to the management group. See “Best Practices for Scheduling Remote Snapshots” on page 36.

## What the System Does

**Best Practice: 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 36.

**If you selected an existing volume to become the remote volume**, the system alerts you that all the data on the existing volume will be deleted, but that a snapshot of all the existing data will be created first. The snapshot that is then created retains all the volume's data.

1. Type a name for that snapshot in the alert.
2. Click Yes 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. Use these steps 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 to bring it to the front.
3. Select the schedule you want.
4. Click Schedule Tasks on the Details tab and 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 to bring it to the front.
3. Select the schedule you want.
4. Click Schedule Tasks on the Details tab and 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

### To Edit the Remote Snapshot Schedule

1. In the navigation window, select the primary volume that has the schedule you want to edit.
2. Click the Schedules tab and select the schedule to edit.
3. Click Schedule Tasks and select Edit Schedule.

   The Edit Schedule Remote Snapshot window opens, shown in Figure 2.16.

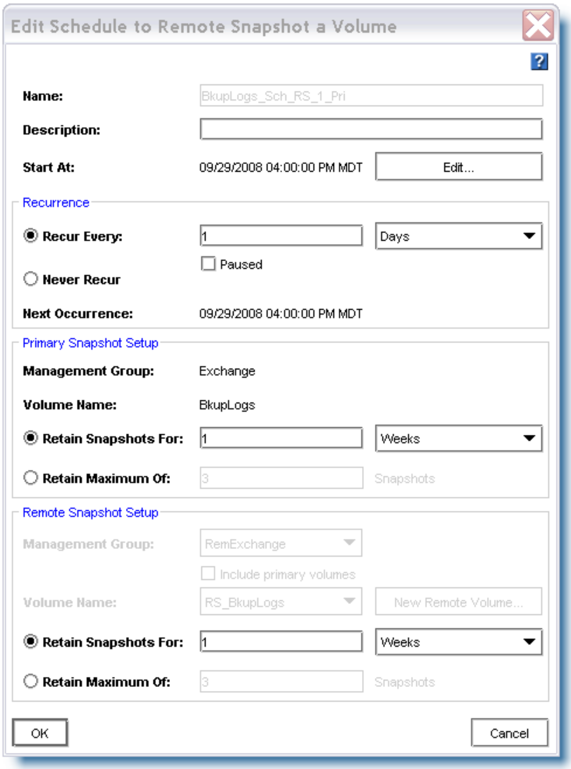

Edit Schedule to Remote Snapshot a Volume

Name: BkupLogs_Sch_RS_1_Pri

Description:

Start At: 09/29/2008 04:00:00 PM MDT Edit...

Recurrence

● Recur Every: 1 Days

□ Paused

○ Never Recur

Next Occurrence: 09/29/2008 04:00:00 PM MDT

Primary Snapshot Setup

Management Group: Exchange

Volume Name: BkupLogs

● Retain Snapshots For: 1 Weeks

○ Retain Maximum Of: 3 Snapshots

Remote Snapshot Setup

Management Group: RemExchange

□ Include primary volumes

Volume Name: RS_BkupLogs New Remote Volume...

● Retain Snapshots For: 1 Weeks

○ Retain Maximum Of: 3 Snapshots

Figure 2.16. Editing the schedule to remote snapshot a volume

4. Change the desired information.
5. Click OK.

## 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 to bring it to the front.
3. Select the schedule you want to delete.
4. Click Schedule Tasks and select Delete Schedule.

   A confirmation message opens.

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)

### 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. Find information about scripting in the *CLIQ — The SAN/iQ® Command-Line Interface User Manual* and in sample scripts available on the LeftHand Networks website.

## 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 promote 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 to use the wizard.

## Making a Primary Volume Into a Remote Volume

Make any primary volume into a remote volume. First, the system takes a snapshot of the primary volume to preserve the existing data that are 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 and select Failover/Failback Volume.

   The Volume Failover/Failback wizard opens, shown in Figure 2.17.

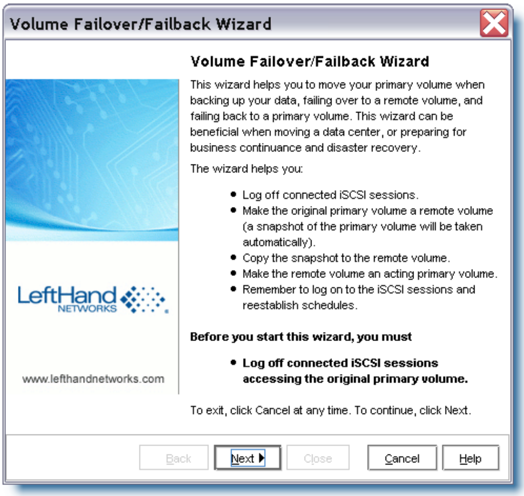

Volume Failover/Failback Wizard

This wizard helps you to move your primary volume when backing up your data, failing over to a remote volume, and failing back to a primary volume. This wizard can be beneficial when moving a data center, or preparing for business continuance and disaster recovery.

The wizard helps you:

- Log off connected iSCSI sessions.
- Make the original primary volume a remote volume (a snapshot of the primary volume will be taken automatically).
- Copy the snapshot to the remote volume.
- Make the remote volume an acting primary volume.
- Remember to log on to the iSCSI sessions and reestablish schedules.

**Before you start this wizard, you must**

- **Log off connected iSCSI sessions accessing the original primary volume.**

To exit, click Cancel at any time. To continue, click Next.

Figure 2.17. Opening the Volume Failover/Failback wizard

3. Click Next.

   You next select the reason you are failing over the volume, shown in Figure 2.18.

- 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 49.

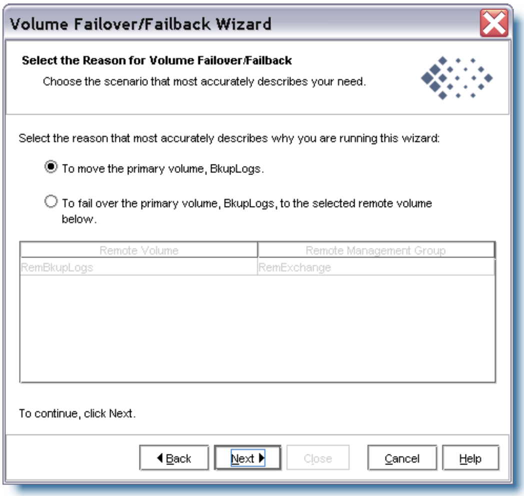

**Select the Reason for Volume Failover/Failback**

Choose the scenario that most accurately describes your need.

Select the reason that most accurately describes why you are running this wizard:

● To move the primary volume, BkupLogs.

○ To fail over the primary volume, BkupLogs, to the selected remote volume below.

| Remote Volume | Remote Management Group |
|---|---|
| RemBkupLogs | RemExchange |

To continue, click Next.

Figure 2.18. Selecting to move the primary volume

4. Select the first choice, to move the volume, and click Next.

   The next window reminds you to disconnect any iSCSI sessions connected to the volume, as shown in Figure 2.19.

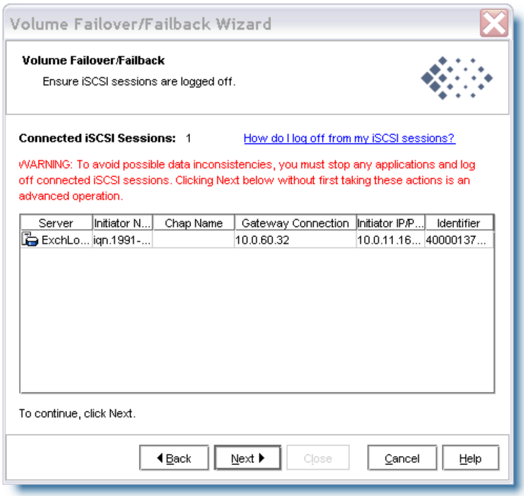

**Volume Failover/Failback**

Ensure iSCSI sessions are logged off.

**Connected iSCSI Sessions:** 1 How do I log off from my iSCSI sessions?

WARNING: To avoid possible data inconsistencies, you must stop any applications and log off connected iSCSI sessions. Clicking Next below without first taking these actions is an advanced operation.

| Server | Initiator N... | Chap Name | Gateway Connection | Initiator IP/P... | Identifier |
|---|---|---|---|---|---|
| ExchLo... | iqn.1991-... | | 10.0.60.32 | 10.0.11.16... | 40000137... |

To continue, click Next.

Figure 2.19. Reminder to disconnect the iSCSI sessions listed

5. Disconnect the iSCSI sessions, if any are displayed, and click Next.

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

6. Next, you can type a name and a description for the snapshot that will be taken of the current primary volume, shown in Figure 2.20.

   This snapshot preserves the existing data on the volume.

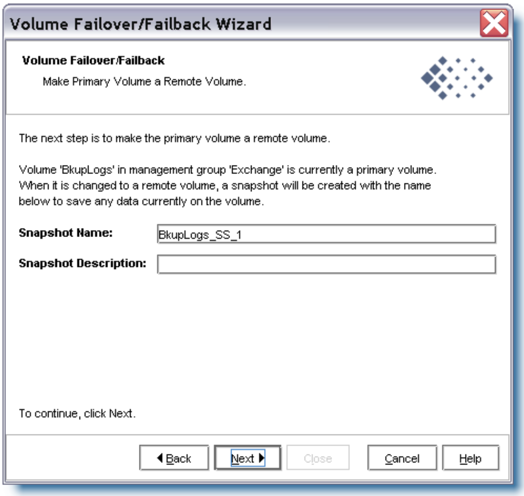
Volume Failover/Failback Wizard
Volume Failover/Failback
Make Primary Volume a Remote Volume.
◀ Back
Next ▶
Close
Cancel
Help

The next step is to make the primary volume a remote volume.

Volume 'BkupLogs' in management group 'Exchange' is currently a primary volume. When it is changed to a remote volume, a snapshot will be created with the name below to save any data currently on the volume.

Snapshot Name: BkupLogs_SS_1

Snapshot Description:

To continue, click Next.

Figure 2.20. Creating a snapshot before making a primary volume into a remote volume

7. Click Next.

8. The next window is where you designate the destination for copying the snapshot to a remote snapshot, shown in Figure 2.21.

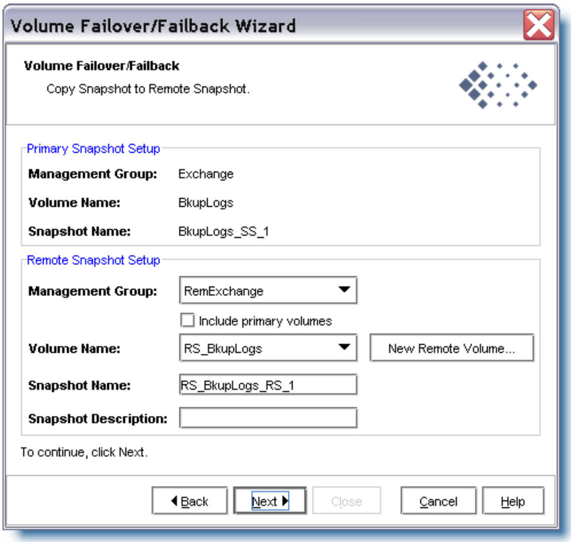

Volume Failover/Failback Wizard

**Volume Failover/Failback**

Copy Snapshot to Remote Snapshot.

Primary Snapshot Setup

| | |
|---|---|
| **Management Group:** | Exchange |
| **Volume Name:** | BkupLogs |
| **Snapshot Name:** | BkupLogs_SS_1 |

Remote Snapshot Setup

| | | |
|---|---|---|
| **Management Group:** | RemExchange | |
| | ☐ Include primary volumes | |
| **Volume Name:** | RS_BkupLogs | New Remote Volume... |
| **Snapshot Name:** | RS_BkupLogs_RS_1 | |
| **Snapshot Description:** | | |

To continue, click Next.

Figure 2.21. Assigning the destination for copying to the remote

9. The final step is to make the remote volume into an acting primary volume, as shown in Figure 2.22.

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

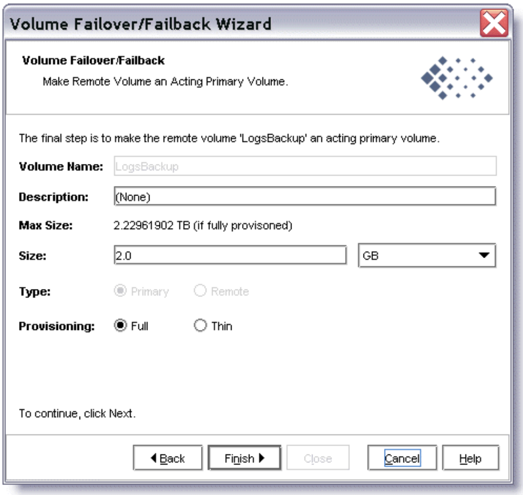

Volume Failover/Failback

Make Remote Volume an Acting Primary Volume.

The final step is to make the remote volume 'LogsBackup' an acting primary volume.

Volume Name: LogsBackup

Description: (None)

Max Size: 2.22961902 TB (if fully provisoned)

Size: 2.0 GB

Type: ● Primary ○ Remote

Provisioning: ● Full ○ Thin

To continue, click Next.

Figure 2.22. Making the remote volume into an acting primary volume

**Note:** You cannot make a remote volume into a primary volume while a remote snapshot is in progress. Wait until the remote snapshot copy is complete before making the remote volume into a primary volume, or cancel the in-progress remote copy.

10. 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, as shown in Figure 2.23.

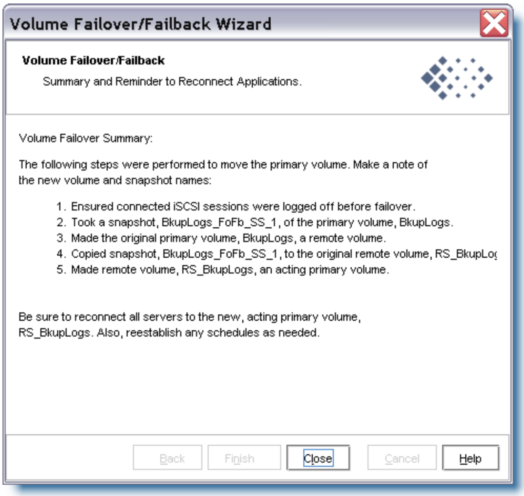

Volume Failover/Failback

Summary and Reminder to Reconnect Applications.

Volume Failover Summary:

The following steps were performed to move the primary volume. Make a note of the new volume and snapshot names:

1. Ensured connected iSCSI sessions were logged off before failover.
2. Took a snapshot, BkupLogs_FoFb_SS_1, of the primary volume, BkupLogs.
3. Made the original primary volume, BkupLogs, a remote volume.
4. Copied snapshot, BkupLogs_FoFb_SS_1, to the original remote volume, RS_BkupLo
5. Made remote volume, RS_BkupLogs, an acting primary volume.

Be sure to reconnect all servers to the new, acting primary volume, RS_BkupLogs. Also, reestablish any schedules as needed.

Figure 2.23. Final summary and reconnect sessions reminder

## 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 and select Failover/Failback Volume.

   The Volume Failover/Failback wizard opens, shown in Figure 2.24.

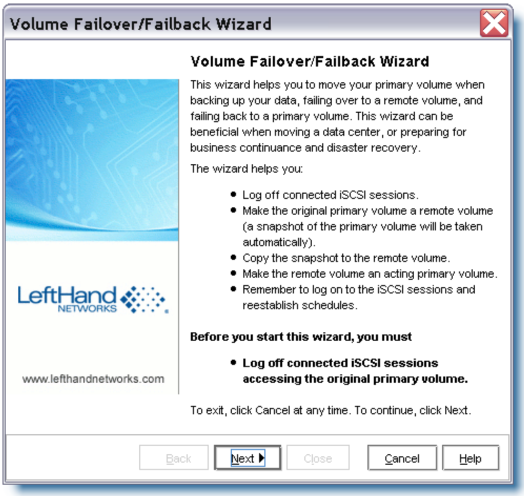

Volume Failover/Failback Wizard

Volume Failover/Failback Wizard

This wizard helps you to move your primary volume when backing up your data, failing over to a remote volume, and failing back to a primary volume. This wizard can be beneficial when moving a data center, or preparing for business continuance and disaster recovery.

The wizard helps you:

- Log off connected iSCSI sessions.
- Make the original primary volume a remote volume (a snapshot of the primary volume will be taken automatically).
- Copy the snapshot to the remote volume.
- Make the remote volume an acting primary volume.
- Remember to log on to the iSCSI sessions and reestablish schedules.

**Before you start this wizard, you must**

- **Log off connected iSCSI sessions accessing the original primary volume.**

To exit, click Cancel at any time. To continue, click Next.

Figure 2.24.  Opening the Volume Failover/Failback wizard

3. Click Next.

   You next select the reason you are failing over the volume, shown in Figure 2.25.

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

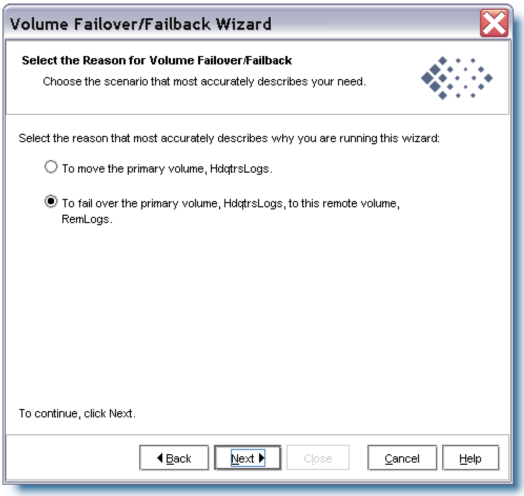

Volume Failover/Failback Wizard

**Select the Reason for Volume Failover/Failback**

Choose the scenario that most accurately describes your need.

Select the reason that most accurately describes why you are running this wizard:

○ To move the primary volume, HdqtrsLogs.

● To fail over the primary volume, HdqtrsLogs, to this remote volume, RemLogs.

To continue, click Next.

Figure 2.25. Selecting to failover to a remote volume

4. Select the second choice, to failover to the remote volume, and click Next.

   The next window reminds you to disconnect any iSCSI sessions connected to the volume, as shown in Figure 2.26.

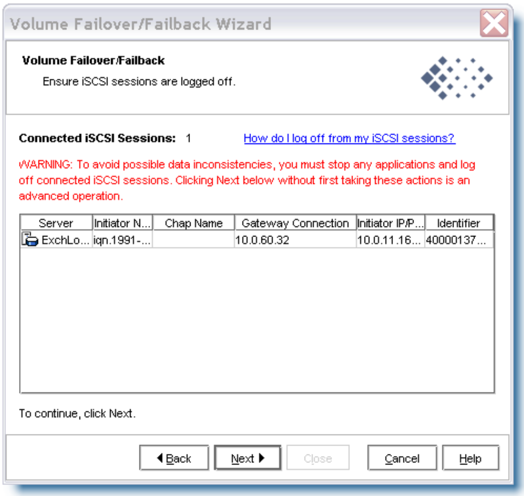

Volume Failover/Failback Wizard

**Volume Failover/Failback**

Ensure iSCSI sessions are logged off.

**Connected iSCSI Sessions:** 1 How do I log off from my iSCSI sessions?

WARNING: To avoid possible data inconsistencies, you must stop any applications and log off connected iSCSI sessions. Clicking Next below without first taking these actions is an advanced operation.

| Server | Initiator N... | Chap Name | Gateway Connection | Initiator IP/P... | Identifier |
|---|---|---|---|---|---|
| ExchLo... | iqn.1991-... |  | 10.0.60.32 | 10.0.11.16... | 40000137... |

To continue, click Next.

Figure 2.26. Reminder to disconnect any iSCSI sessions listed

5. The final step is to make the remote volume into an acting primary volume, as shown in Figure 2.27.

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

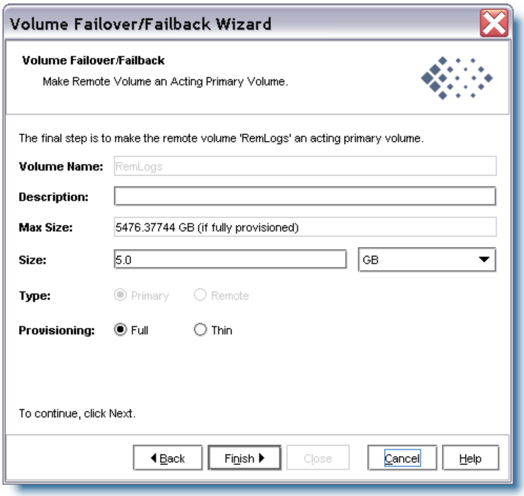

Volume Failover/Failback Wizard

**Volume Failover/Failback**

Make Remote Volume an Acting Primary Volume.

The final step is to make the remote volume 'RemLogs' an acting primary volume.

**Volume Name:** RemLogs

**Description:**

**Max Size:** 5476.37744 GB (if fully provisioned)

**Size:** 5.0 GB

**Type:** ● Primary ○ Remote

**Provisioning:** ● Full ○ Thin

To continue, click Next.

Figure 2.27. Making the remote volume into an acting primary volume

6. Click Finish.

The final window of the wizard displays a summary of the actions and a reminder to reconnect your iSCSI sessions, as shown in Figure 2.28.

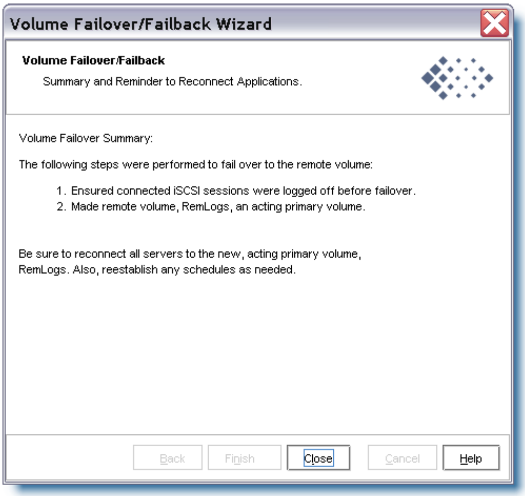

Volume Failover/Failback

Summary and Reminder to Reconnect Applications.

Volume Failover Summary:

The following steps were performed to fail over to the remote volume:

1. Ensured connected iSCSI sessions were logged off before failover.
2. Made remote volume, RemLogs, an acting primary volume.

Be sure to reconnect all servers to the new, acting primary volume, RemLogs. Also, reestablish any schedules as needed.

Figure 2.28. Final summary and reconnect sessions reminder

# Resuming Production After Failover

After failover occurs, three scenarios exist for resuming production.

- Failback returns operations to the original primary site once it is restored.
- Make the backup site into 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 will usually be 2 snapshots or volumes that have conflicting data. Recovering and synchronizing such data depends on multiple factors, including the application involved. For more detail about synchronizing, see Table 1.1, “Remote Copy Glossary,” on page 2.

### Example Scenario

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

## Time Line of Failover

Table 2.2. Time line of failover

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

## Data that now needs to be synchronized

- 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

Once 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

**1. Create Snapshots of Data**

Create snapshots that contain the data that you need to synchronize. The steps to create those snapshots are described in Table 2.3.

Table 2.3. 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. | OrigPrimaryVol_1<br>OrigPrimarySS_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. | | Remotevol_1<br>ActPrimarySS_1 | Returns the remote management group to its original configuration. In addition, you capture the 1:33 to 2:00 p.m. data. |

**2. Synchronize the Data**

"Synchronize" the snapshots OrigPrimarySS_1 and ActPrimarySS_1 that were created in Steps 2 and 3 of Table 2.3.

- In the simplest case, to synchronize the snapshots, remote copy the remote snapshot back to the original primary volume. For more detail about synchronizing, see Table 1.1, "Remote Copy Glossary," on page 2.

## 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 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.

   As your target, create a remote volume, which will later be converted into your primary volume.
3. Convert the remote volume into a primary volume.
4. Make the acting primary volume into the remote volume.

   This creates a snapshot of that volume.
5. Configure a new timetable for the scheduled remote snapshots on the new primary volume.
6. 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 require that you delete any snapshots that were created after the snapshot that is rolled back to.

**New for Release 8.0**

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. See detailed information in Chapter 14 — SmartClone volumes in the *LeftHand SAN User Manual*.

## 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 than the original.

**Prerequisites**

- Stop applications from accessing the volume.

**Warning:** 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 select Roll Back Volume.

   The Roll Back Volume window opens, shown in Figure 2.29 and Figure 2.30.

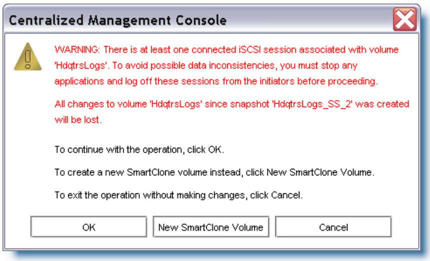

Centralized Management Console

WARNING: There is at least one connected iSCSI session associated with volume 'HdqtrsLogs'. To avoid possible data inconsistencies, you must stop any applications and log off these sessions from the initiators before proceeding.

All changes to volume 'HdqtrsLogs' since snapshot 'HdqtrsLogs_SS_2' was created will be lost.

To continue with the operation, click OK.

To create a new SmartClone volume instead, click New SmartClone Volume.

To exit the operation without making changes, click Cancel.

Figure 2.29. Rolling back a primary volume if iSCSI sessions are still connected

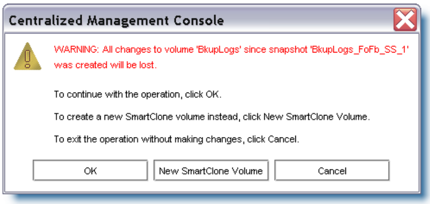

Centralized Management Console

WARNING: All changes to volume 'BkupLogs' since snapshot 'BkupLogs_FoFb_SS_1' was created will be lost.

To continue with the operation, click OK.

To create a new SmartClone volume instead, click New SmartClone Volume.

To exit the operation without making changes, click Cancel.

Figure 2.30. 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. To roll back a remote volume, make the remote volume into a primary volume and perform the steps previously mentioned.

# 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. See detailed information in Chapter 14 — SmartClone volumes in the *LeftHand SAN User Manual*.

## Creating a Split Mirror

To create a split mirror, perform these general steps:

- 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**

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 select Edit Management Group.

   The Edit Management Groups window opens, shown in Figure 2.31.

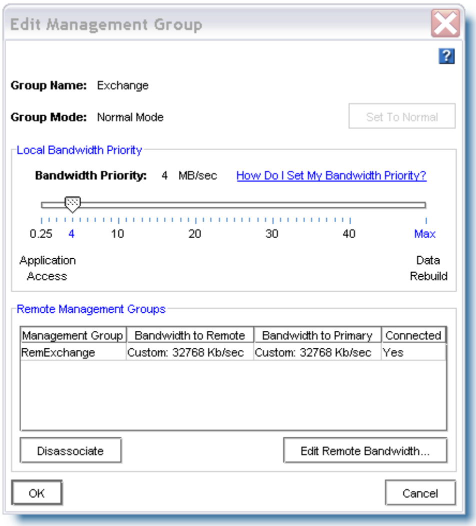

Edit Management Group

Group Name: Exchange

Group Mode: Normal Mode

Local Bandwidth Priority

Bandwidth Priority: 4 MB/sec

Remote Management Groups

| Management Group | Bandwidth to Remote | Bandwidth to Primary | Connected |
|---|---|---|---|
| RemExchange | Custom: 32768 Kb/sec | Custom: 32768 Kb/sec | Yes |

Figure 2.31. Editing a 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.

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

6. 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.

7. Click OK to return to the navigation window.

# Sample Remote Copy Configurations 3

## Overview

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

### Achieving High Availability

Creating remote snapshots in remote locations with Remote Copy can ensure 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 volume replication within a cluster, ensures 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 3.1 illustrates this simple high availability configuration.

- Configure primary and backup application servers in both the primary and backup locations.

  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.

## Configuration Diagram

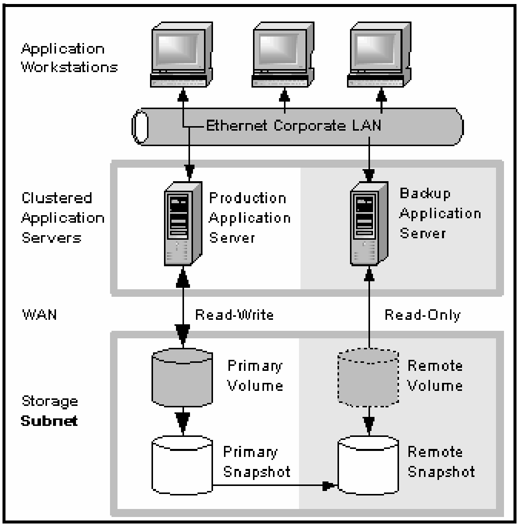

Figure 3.1. High availability example configuration

## 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. As shown in Figure 3.2, 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.

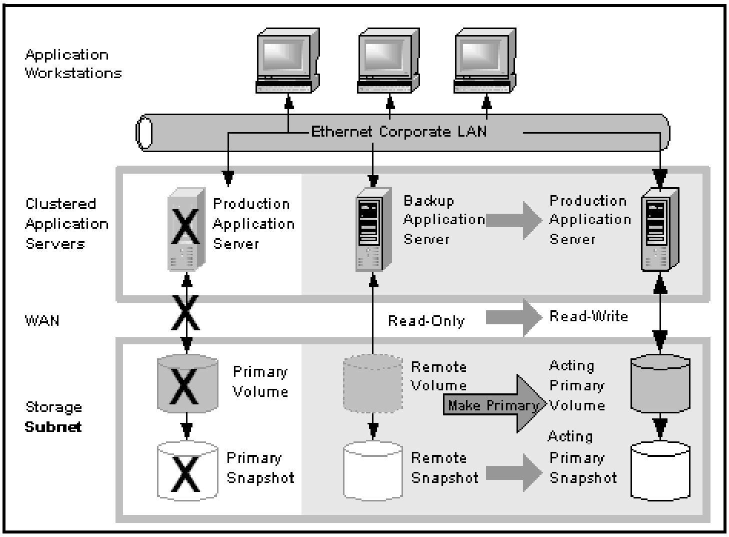

Figure 3.2. High availability configuration during failover

### 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 will be 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. Find information about scripting in the *CLIQ — The SAN/iQ® Command-Line Interface User Manual* and in sample scripts available on the LeftHand Networks website.
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 two failback options:

- Resume operations using the original production server, and return the backup volumes to their original remote status, as illustrated in Figure 3.3. This will require migration back onto the production volumes of data that were written to the backup volumes since the failure.
- 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. See "Synchronizing the Data Between the Acting Primary Volume and the Original Primary Volume" on page 55.

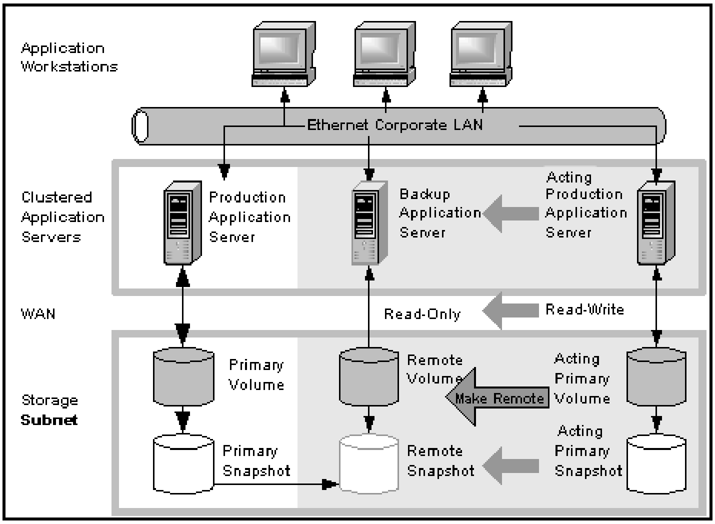

Figure 3.3. High availability configuration during failback

# Best Practices

## Remote Snapshots with Volume Replication

Use remote snapshots in conjunction with local, synchronous volume replication. 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 synchronous volume replication. Volume replication allows you to create up to 4 copies of a volume on the same cluster of storage nodes as the primary volume. The only limitation is that the cluster must contain at least as many storage nodes as replicas of the volume. Replicating the volume within the cluster ensures that if a storage node in the cluster goes down, replicas of the volume elsewhere in the cluster will still be available. For detailed information about volume replication, see the chapter “Provisioning Storage” in the *LeftHand SAN User Manual* for details.

## Example Configuration

This example, illustrated in Figure 3.4, uses 3 storage nodes per cluster. However, this scenario can use any number of storage nodes. Information about creating clusters and volumes can be found in the *LeftHand SAN User Manual*.

- In the production location, create a management group and a cluster of 3 storage nodes.
- Create volumes on the cluster, and set the replication level to 2-Way.
- 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 3 storage nodes.
- Create a schedule for making remote snapshots of the primary volume. See “Scheduling Remote Snapshots” on page 36.

**Note:** Volume replication levels are set independently for primary and remote volumes.

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

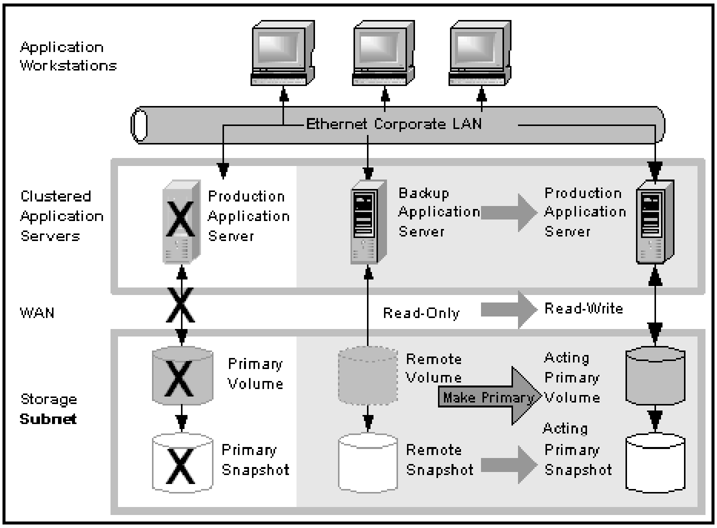

Figure 3.4. High Availability During Failover - Example Configuration

## 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 nodes 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 3.5 shows an example configuration for disaster recovery.

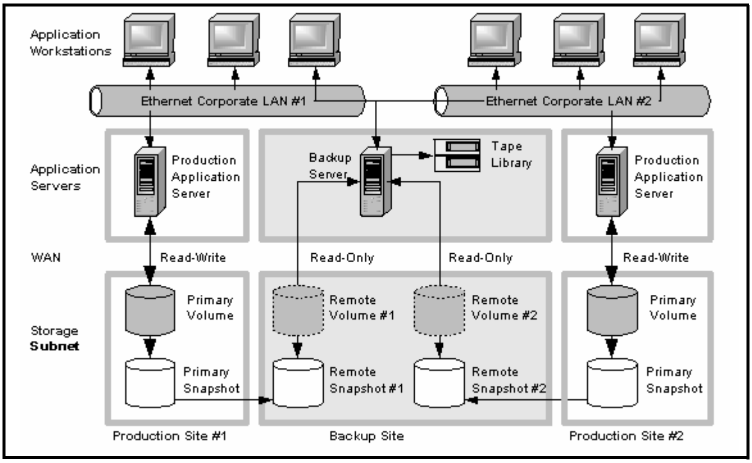

Figure 3.5. Affordable disaster recovery example configuration

## How This Works for Affordable Disaster Recovery

If the storage nodes 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 3.6. If you created tape backups, you can recover data from tape backups, as shown in Figure 3.7.
- 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 this 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.

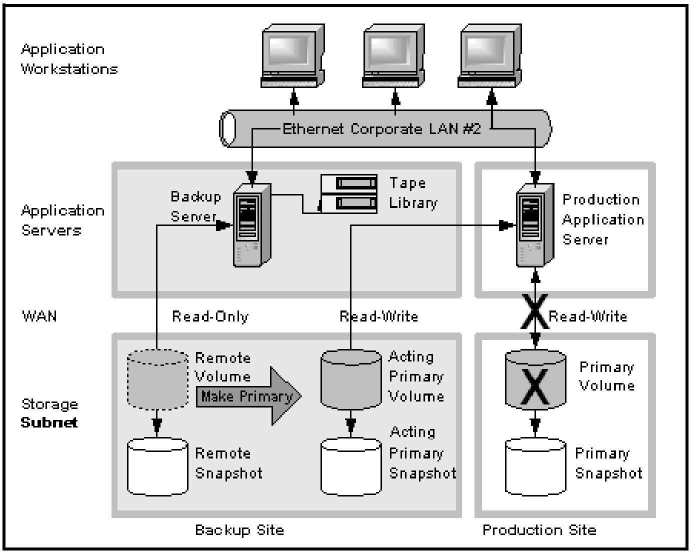

Figure 3.6. Restoring from a remote volume

In Figure 3.6, note the volume labeled Primary Snapshot in the gray area on the left. It originated as a read only back up, but is brought into use as an acting primary.

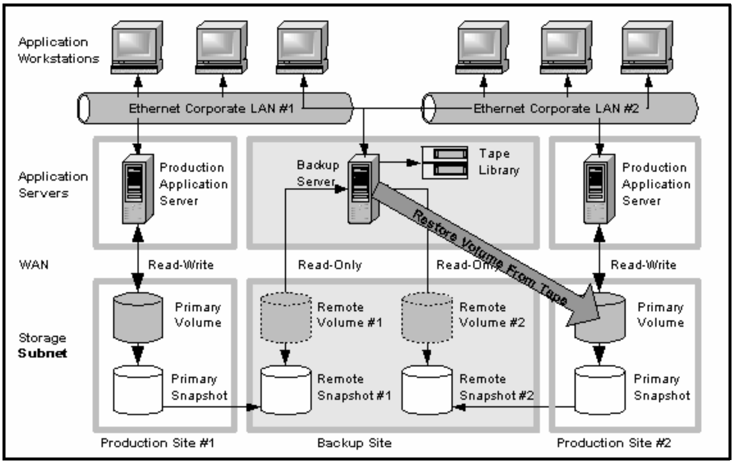

Figure 3.7. Restoring from tape backup

## 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 4 hours to ensure that no more than 4 hours worth of data are lost in an emergency. In addition, you want to retain 1 week's worth of remote snapshots. Retaining 4-hour snapshots for 1 week can result in the accumulation of over 40 remote snapshots. Another approach would be to create 2 remote snapshot schedules for the volume:

- One schedule to create remote snapshots every 4 hours, but only retain the most recent 6 remote snapshots. This will ensure that you do not lose more than 4 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 synchronous replication of the volume within the cluster of storage nodes at the primary geographic site. With synchronous replication, a single storage node can be off-line, and your primary volume will remain intact.

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

### Example Configuration

- In the production location, create a cluster of 3 storage nodes, all with managers.
- Create volumes on the cluster, and set the replication level to 2-Way.
- Create a schedule for making remote snapshots of the primary volume. Set the recurrence to every 4 hours, and retention of remote snapshots to 2 days.

**Note:** You can use the same volume replication configuration on the remote volume as well. However, this replication is configured independently of the volume replication that is configured on the primary volume.

If one of the storage nodes in the primary location fails, the primary volume will still be available. If all of the storage nodes 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 backup server (either local or remote), 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.

See the example configuration illustrated in Figure 3.8.

## Configuration Diagram

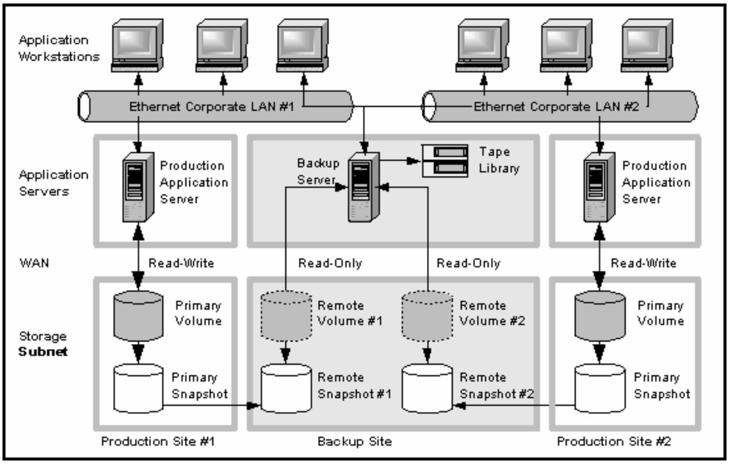

Figure 3.8. Off-site backup and recovery example configuration

# 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 going to 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, you specify a number of primary and remote snapshots that you want to retain. You can retain primary snapshots to facilitate easy rollback of the primary volume. (Retention of snapshots will affect the amount of space that is used in the cluster of storage nodes, so balance the number of snapshots to retain with the amount of space you are willing to use. To roll back to a snapshot that you did not retain, you can still access remote snapshots or tape backups.)

- 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 these data without going to the backup tape.

### Example Configuration

- Retain 3 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 1 week, go to the backup tape.

## Achieving Non-Destructive Rollback

As discussed in “Rolling Back a Primary Volume” on page 58, 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 will need to 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 2 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 Non-Destructive Rollback

To use remote snapshots for non-destructive 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 will remain intact on the other volume.

See Figure 3.9 for an illustration of this configuration.

### Configuration Diagram

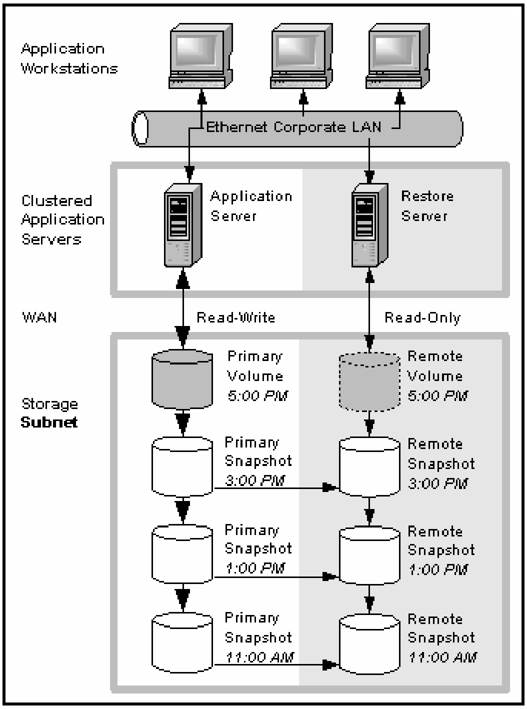

Figure 3.9. Non-destructive rollback example

## How This Configuration Works for Non-Destructive 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 (as 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 3.10 shows rollback of the primary snapshot while leaving the remote snapshots intact.

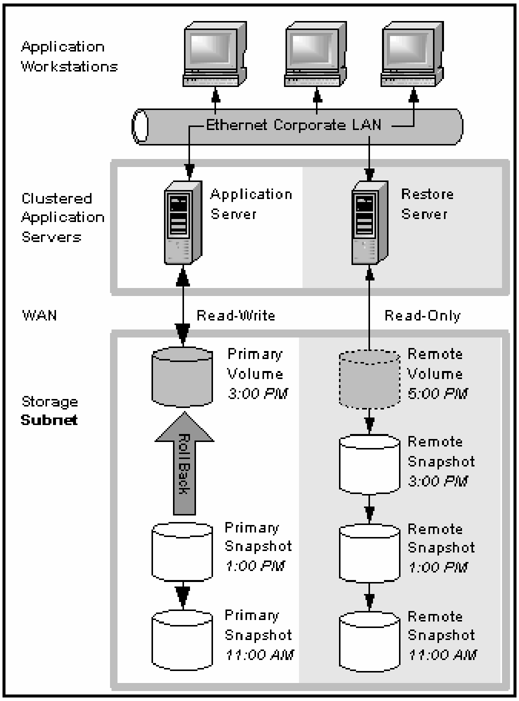

Figure 3.10. Non-destructive rollback from the primary snapshot

- To roll back the remote snapshot, you must first make the remote volume into a primary volume. This will stop scheduled creation of remote snapshots, which may jeopardize your high availability, disaster recovery, or routine backup strategies. Figure 3.11 shows rollback of the remote snapshot.

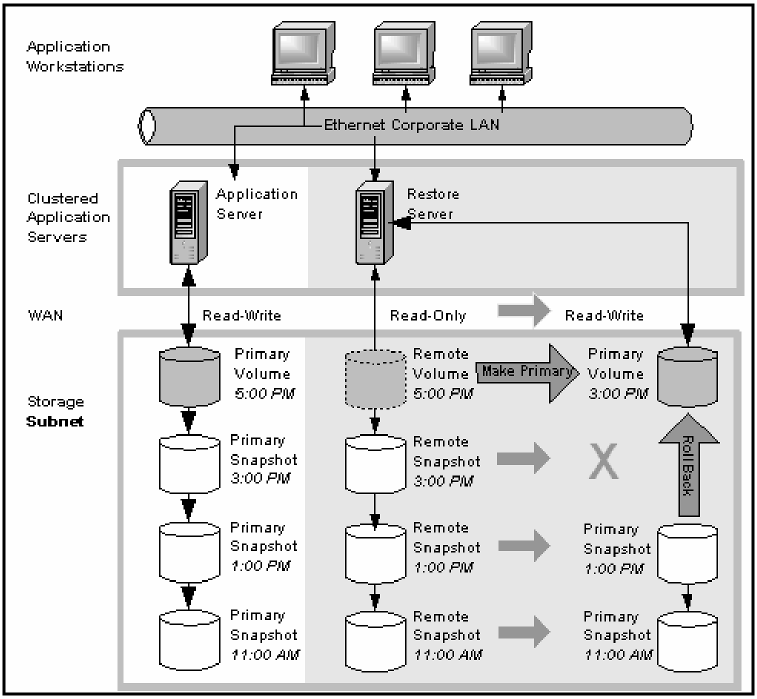

Figure 3.11. Non-destructive rollback from the remote snapshot

## 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 nodes in the remote location with enough space to accommodate the volume. See the example illustrated in Figure 3.12.

## Configuration Diagram

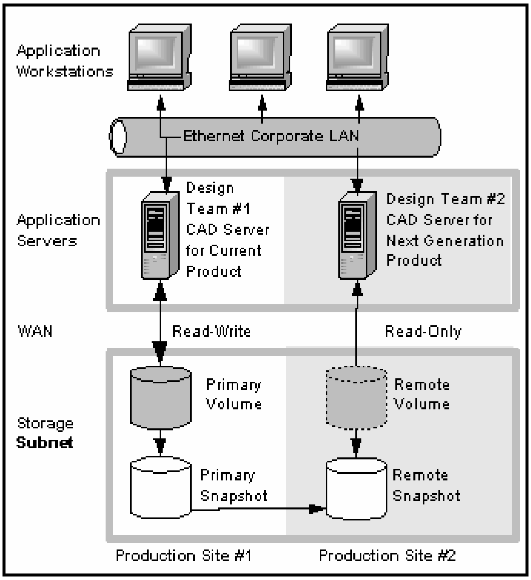

Figure 3.12. Data migration example configuration

## 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 nodes 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 3.13 shows migration of data by making a remote volume into a primary volume.

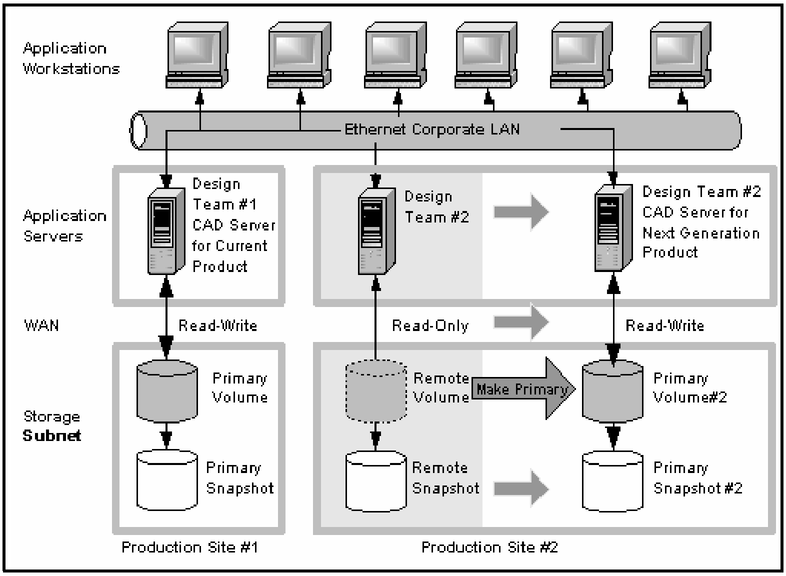

Figure 3.13. Configuration after data migration

# Index